Teradata Active EDW 6650 & 6680

Jim Dietz
Enterprise Product Marketing
4/8/11

Platform Introduction & Overview

TERADATA
Raising Intelligence

Agenda

- New Teradata Active EDW platforms
- Teradata Active EDW 6680
- Teradata Active EDW 6650
- Teradata SSD technology & performance
- Active EDW platform solutions

Teradata **Purpose-Built** Platform Family

NEW Platform Models Released @ PARTNERS October, 2010





	Data Mart Appliance	Extreme Data Appliance	Data Warehouse Appliance	Extreme Performance Appliance	Active Enterprise Data Warehouse
Purpose	Test / Development or Smaller Data Marts	Analytics on Extreme Data Volumes from New Data Types	Data Warehouse or Departmental Data Marts	Extreme Performance for Operational Analytics	Enterprise Scale for both Strategic and Operational Intelligence EDW / ADW
Scalability	Up to 11TB	Up to 186PB	Up to 343TB	Up to 18TB	Up to 92PB
Sub Segment	Departmental Analytics, Entry level EDW	Analytical Archive, Deep Dive Analytics	Strategic Intelligence, Decision Support System, Fast Scan	Operational Intelligence, Lower Volume, High Performance	Active Workloads, Real Time Update, Tactical and Strategic response times
Current Model #	560	1600 1650 (Q1)	2650	4600	5650

TERADATA
Raising Intelligence

Teradata **Purpose-Built** Platform Family





	Data Mart Appliance	Extreme Data Appliance	Data Warehouse Appliance	Extreme Performance Appliance	Active Enterprise Data Warehouse
Purpose	Test / Development or Smaller Data Marts	Analytics on Extreme Data Volumes from New Data Types	Data Warehouse or Departmental Data Marts	Extreme Performance for Operational Analytics	Enterprise Scale for both Strategic and Operational Intelligence EDW / ADW
Scalability	Up to 11TB	Up to 186PB	Up to 343TB	Up to 18TB	Up to 92PB
Sub Segment	Departmental Analytics, Entry level EDW	Analytical Archive, Deep Dive Analytics	Strategic Intelligence Decis Sys	Operational , ce	Active Workloads, Real Time Update, Tactical and Strategic response times
Current Model #	560	1600 1650 (Q1)	2650	4600	6650 / 6680

Public Release 4/11/11

TERADATA. Raising Intelligence

Active EDW – Responding to the Market

The Need for Speed

1. Effective use of hot data is crucial to business
   - > Near real time operational
2. Massive growth in use of advanced analytics
   - > High demands on system resources
3. Reduce energy and floor space while delivering more performance
   - > Data centers are limited
4. Big Data demands flexible solutions for growth in both
   - > Performance
   - > Data capacity

Hot Data Use Case Examples

New Opportunities for Data Warehousing

- BI dashboards
  - > 7 X 24 real time operations
- Financial peak periods
  - > Month end, quarter end
- Cyber Security
  - > Short and long term threats
- Operational reporting
  - > Claims processed
  - > Inventory instant status
- Machine generated data
  - > Rapid response
  - > Fast analytics

SSD & Mixed Storage Potential Applications

Steep Growth in Advanced Analytics

- According to a 2009 survey*
  - > 38% practicing advanced analytics
  - > 85% practicing it by 2012
- Advanced analytics consumes processing horsepower and information
  - > For data preparation, scoring, and regression analysis
  - > Chew up a huge slice of CPU and I/O bandwidth.
- Two forms of advanced analytics
  - > *query-based analytics* (complex SQL define recent business events)
  - > *predictive analytics* (data mining and statistical methods anticipate future events).

*TDWI

Both drive the need for performance on data

TERADATA
Raising Intelligence

Teradata Active EDW 6680 & 6650

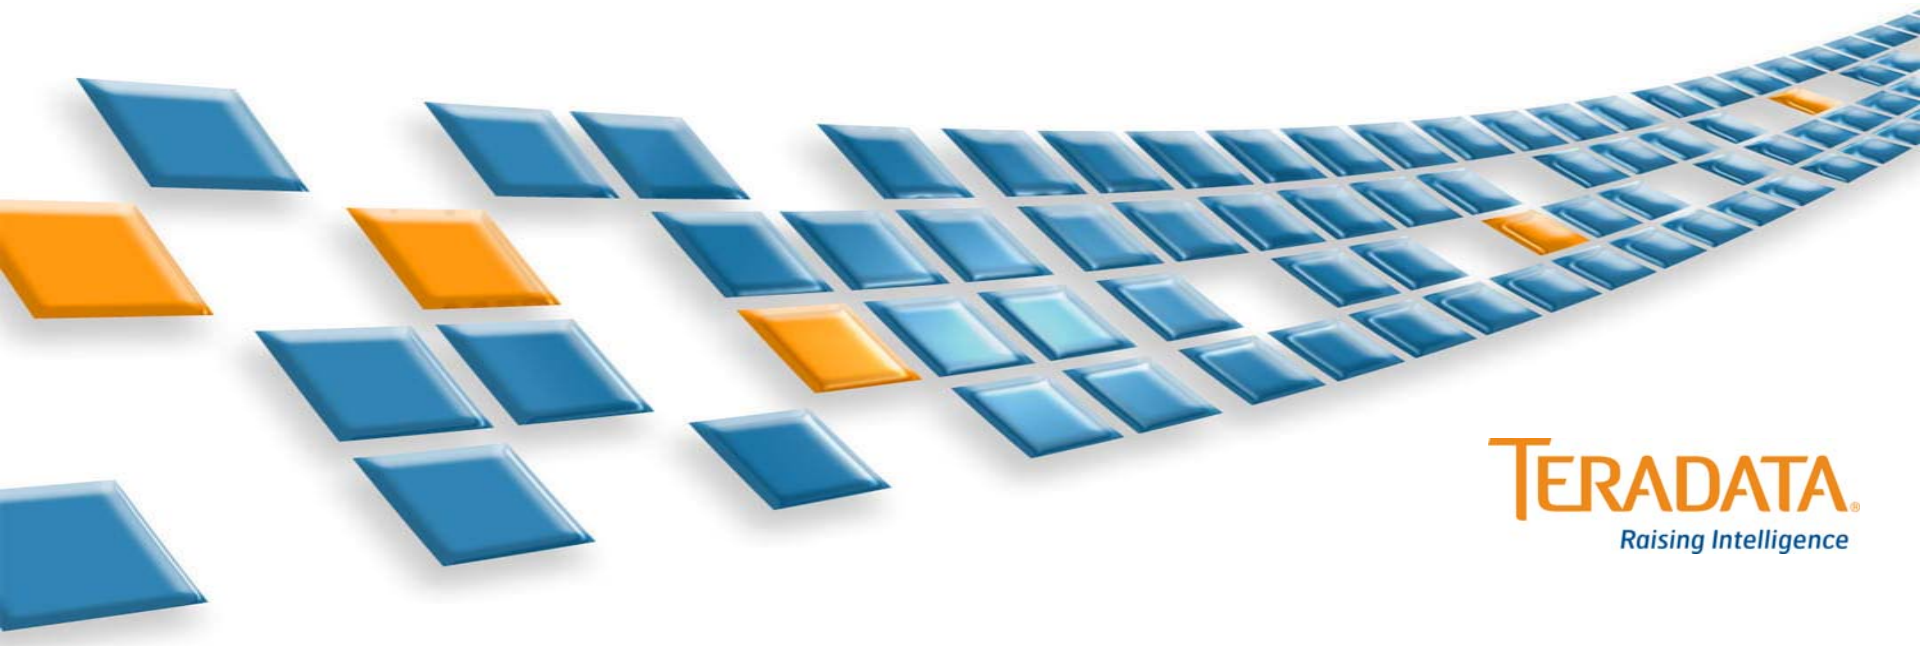

Teradata Active EDW 6680: Efficient Performance

- What's new?
  - > Mixed storage for data warehouse
    - SSD for I/O performance
    - HDD for capacity
  - > Extend Teradata Virtual Storage (TVS) for different drive types
    - Provides automatic data migration
    - Hot data in SSDs & warm/cold data in HDD
  - > Single cabinet for each 2 node clique
    - Efficient performance per KW & footprint
- What's the same as 5650?
  - > Nodes, HDD storage, BYNET, etc.
  - > Coexistence (with future generations)

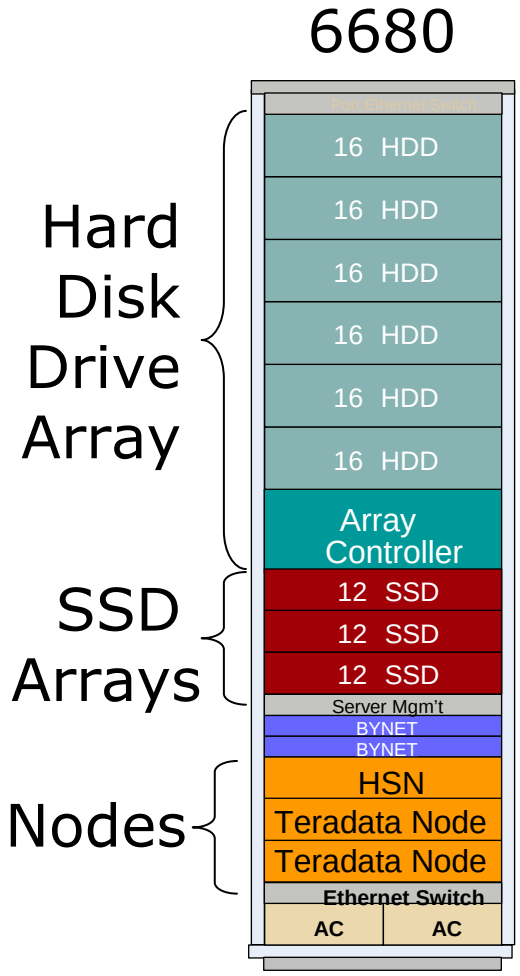

TERADATA. Raising Intelligence

Teradata Active Enterprise Data Warehouse 6680

Nodes	Based on Intel six-core Xeon processors
Storage - Mixed	SSD: 300GB Enterprise Flash Drive HDD: 300GB, 450GB, 600GB drives
Configuration Flexibility	4 configurations: 12 to 18 SSD per node and 36 to 48 HDD per node
Data capacity	From 3.8TB Up to 20PB
Scalability	Scales up to 4,096 nodes
Availability features	RAID 1, automatic node failover and recovery, performance continuity with hot standby nodes, fallback, BAR, Dual Active systems
Software	Teradata 13.10, TVS, 64-bit SUSE Linux
Workload management	Full Teradata Active System Manager
System management	Single operational view of System and Viewpoint
Interconnect	Teradata BYNET® V4

= New

TERADATA
Raising Intelligence

Teradata Active EDW 6650: SSD Ready & "Green"

- What's new?
  - > Solid State Drive (SSD) ready
    - – SSD transition platform for investment protection
    - – Optional add in late 2011
  - > Nodes and storage in same cabinet
    - – 1 to 3 nodes + Hot Standby Node in 2 cabinet clique*
    - – Flexible configurations save space and energy costs
- What's the same?
  - > All basic elements of previous 5650 model
    - – Nodes, BYNET®, software, etc
    - – Storage drives and array
  - > Coexistence back to 5500 generation

*Clique is the basic Teradata redundancy grouping of nodes

Active EDW 6650

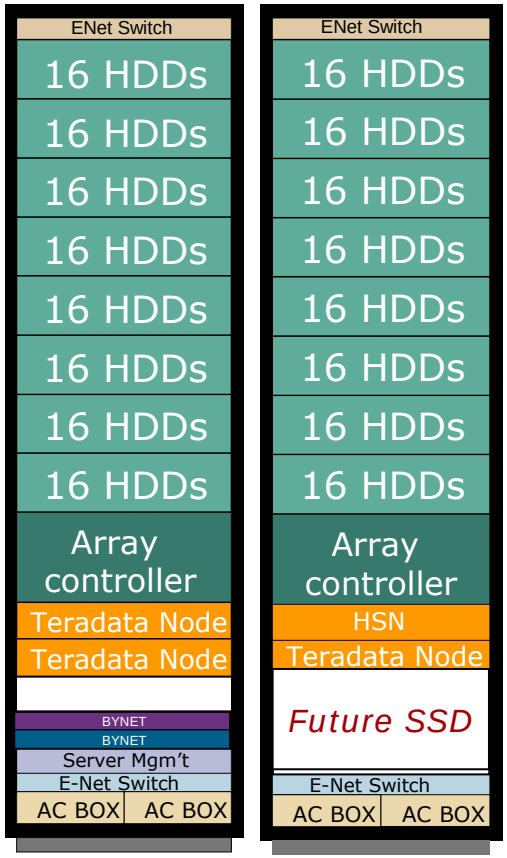

Two cabinet clique
(Full configuration shown)

Teradata Active Enterprise Data Warehouse 6650

Nodes	Based on Intel six-core Xeon processors
Storage	300GB or 450GB HDD enterprise-class drives; SSD in future
Configuration Flexibility	42 HDD per node to 232 HDD per node
Data capacity	From 7.5TB Up to 92PB
Scalability	Scales up to 4,096 nodes
Availability features	RAID 1, automatic node failover and recovery, performance continuity with hot standby nodes, fallback, BAR, Dual Active
Software	Teradata 12 & higher, SUSE Linux
Workload management	Full Teradata Active System Manager
System management	Single operational view and Viewpoint
Interconnect	Teradata BYNET® V4

= New

TERADATA
Raising Intelligence

6650 Upgrade Path to SSD

- Determine data temperature profile
  - > Hot data determines SSD size
  - > Use Teradata Temperature Assessment tool
- Upgrade current 6650 to SSD when performance is needed
  - > Add SSD storage arrays
  - > Add CPU capability
  - > Upgrade to Teradata 13.10 and TVS
  - > Restart system
- Performs as an SSD mixed storage system

Teradata Active EDW Investment Protection

- 6650 supports up to 6 generations of coexistence
  - > Enables full performance from all generations of nodes in the system

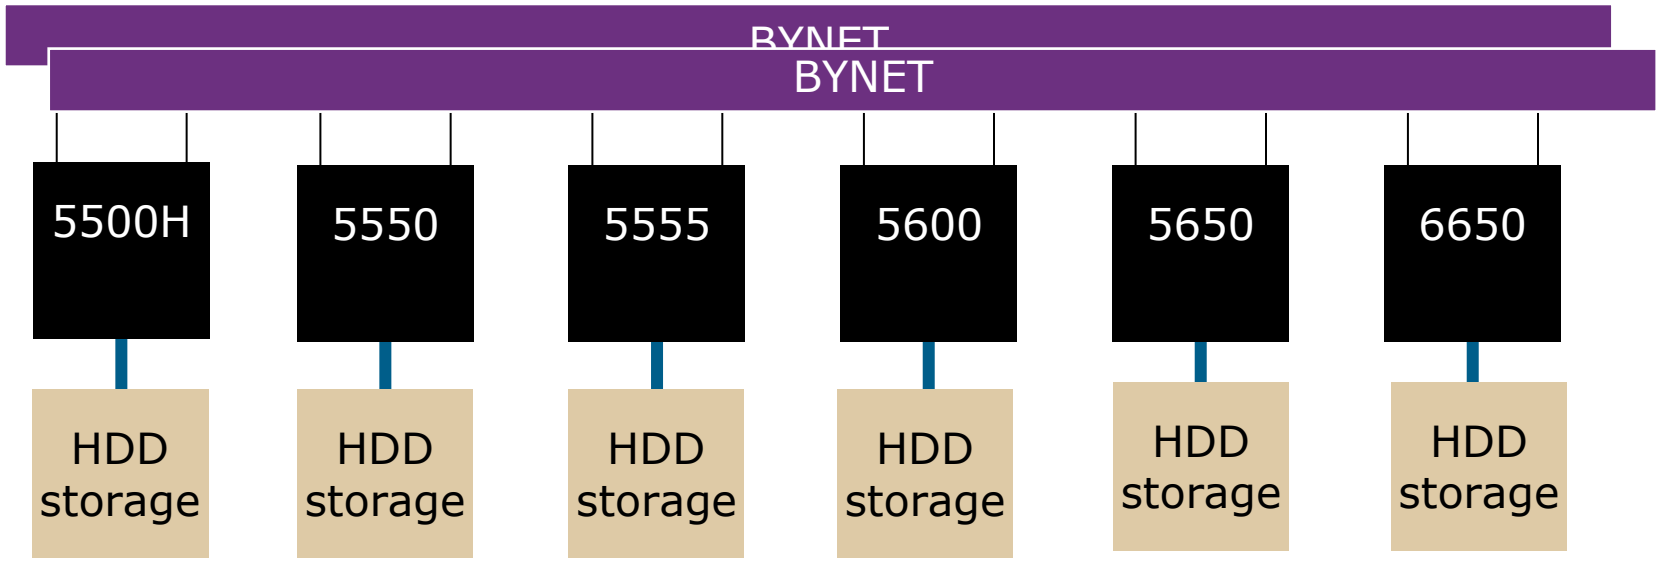

Mixed Storage:
- SSD
- Teradata Virtual Storage
- Performance Impacts

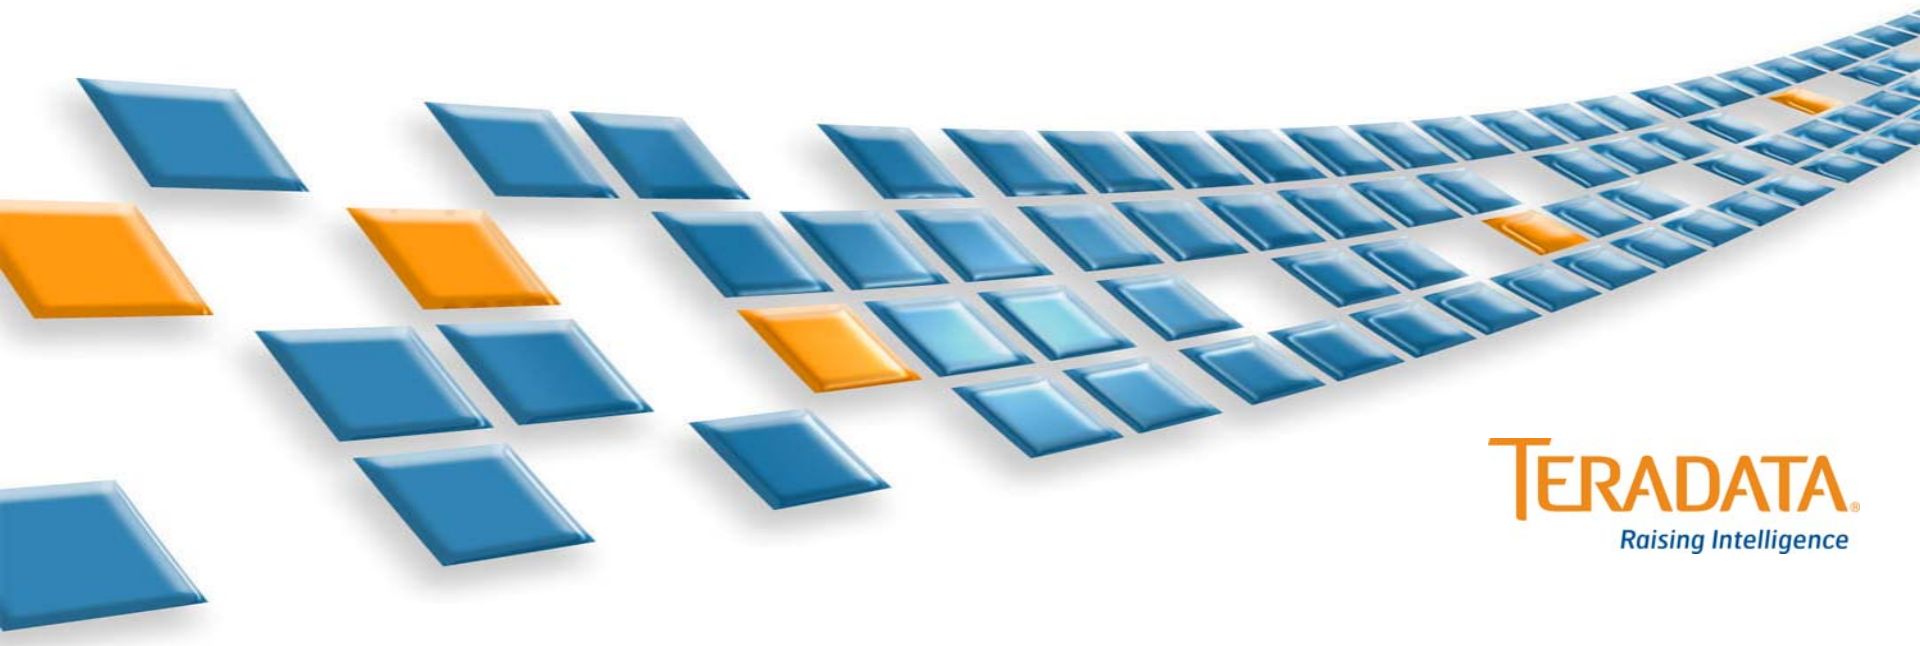

TERADATA
Raising Intelligence

SSD: Drive Performance Catches Up to CPU

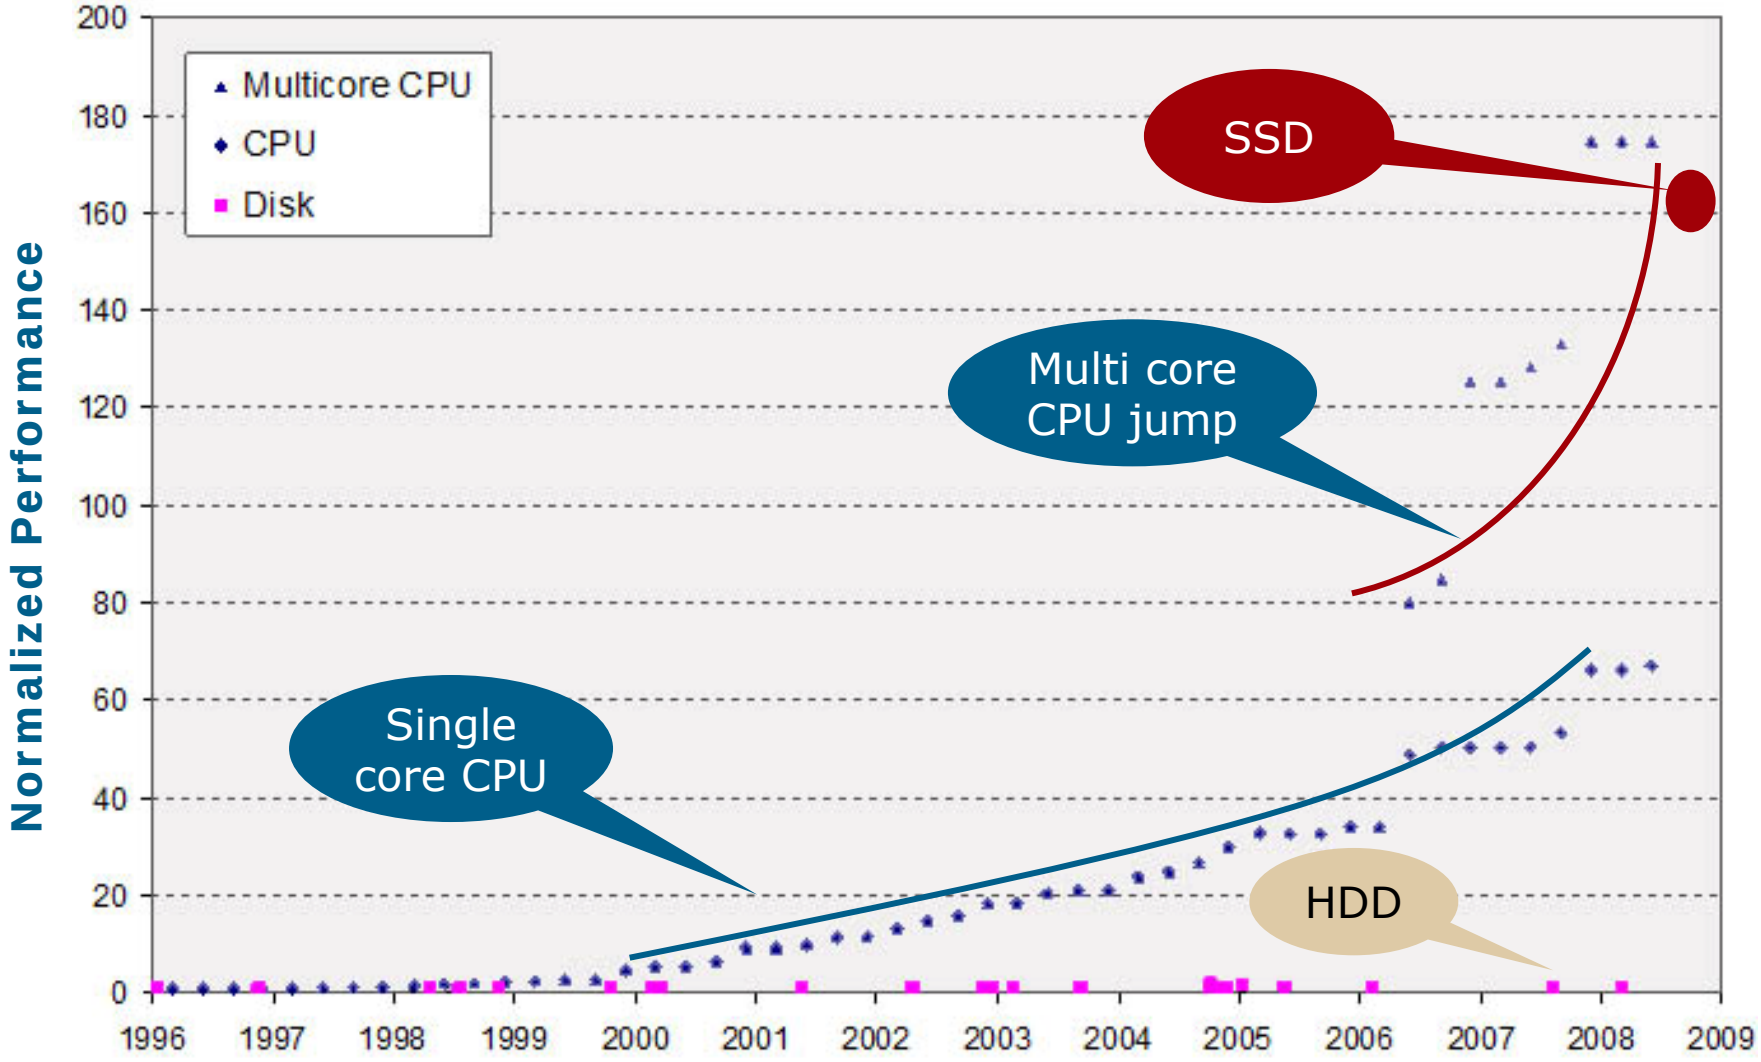

Intel CPUs over Time

TERADATA
Raising Intelligence

One SSD Provides Throughput of 22 HDDs

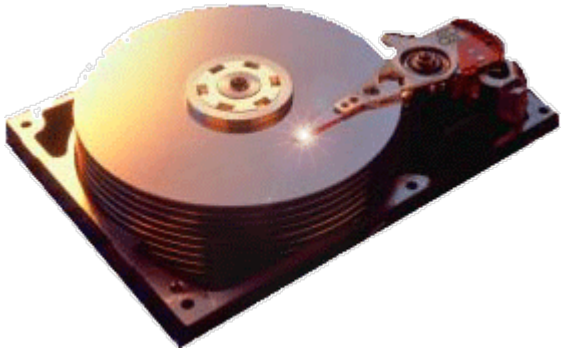

The Teradata "Sweet Spot" for drive use

	Pliant Enterprise SSD	Ratio	Enterprise 15K HDD
Random 80/20 (MB/sec)	**>450**	**22X**	**20**
Random I/O latency (Sec)	10^{-6}	>1000	10^{-3}
I/O per second (4 KB)	10^{5}	150X	10^{2}
Sequential read BW (MB/s)	>450	3X	>150

Where other DW use drives

Today's HDD Based Data Warehouse

Challenge to Balance CPU and Storage I/O

- Ever growing TB per node
  - > Higher drive counts
  - > Larger drive sizes reduce I/O per TB ratio
    - 300 and 450 GB
  - > Performance per data space has eroded
    - 75% lower over 4 years
- How does the Market Respond?
  - > Partially fill disk space – "short stroke"
    - Boosts effective performance per data space
  - > Commercial concessions

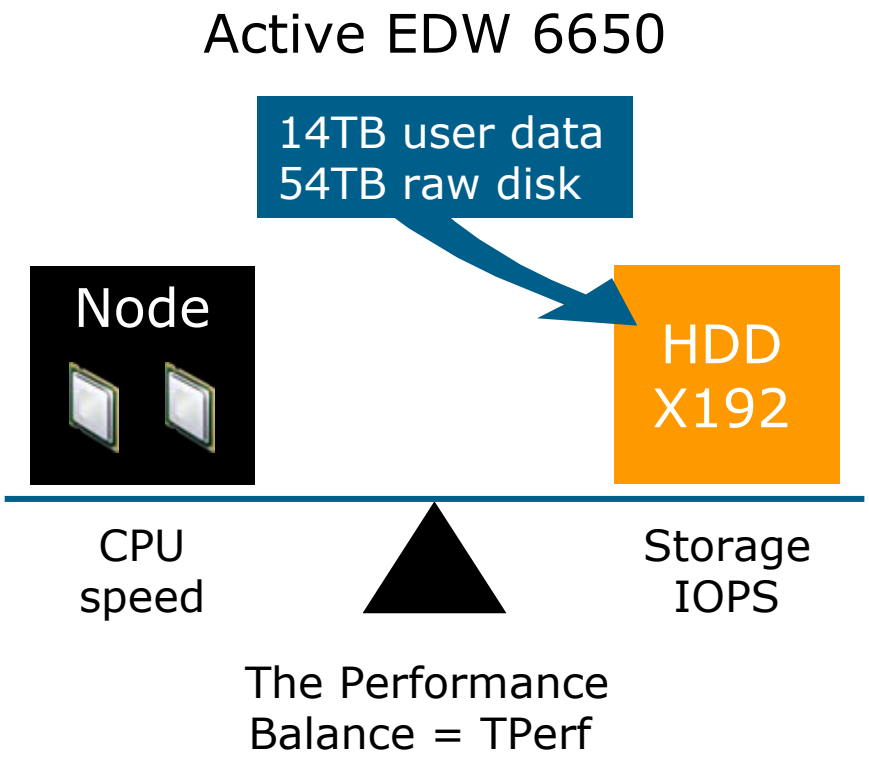

HDD = Hard Disk Drive
All drives = 300GB
IOPS = I/O Operations per Second

Teradata Active EDW with Mixed Storage

- SSD reduces # of drives required
  - > Maintains CPU to IOPS balance
  - > SSD speed = up to 22X HDD for Teradata I/O usage
- Higher performance per TB
  - > Reduced TB per node
  - > Maintains same Node TPerf
  - > Can boost performance per data space 2X – 4X
- Deliver performance without excess data capacity

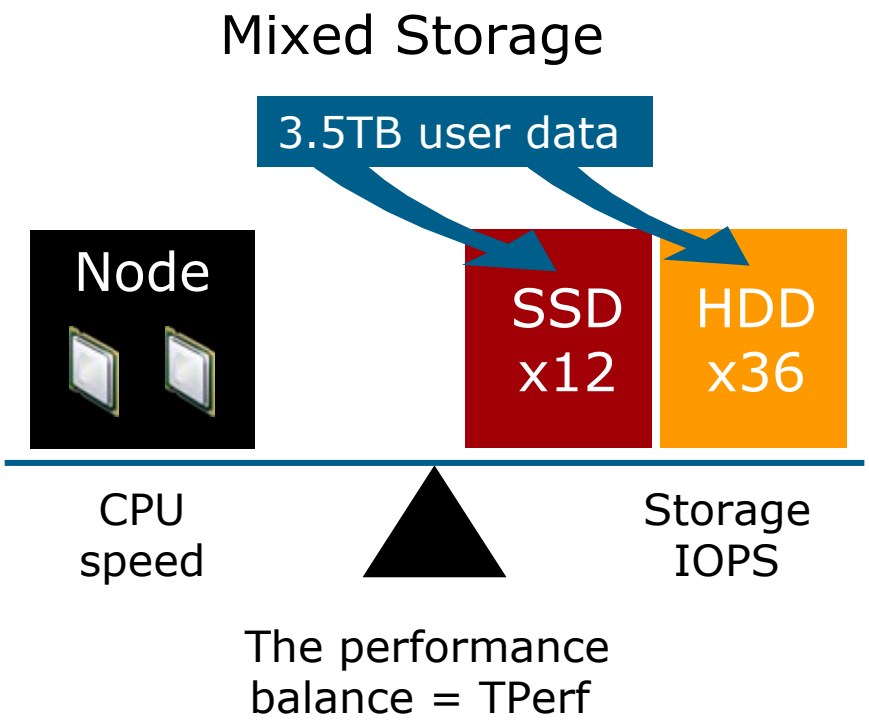

Example configuration:
SSD = Solid State Drive
HDD = Hard Disk Drive
All drives = 300GB
IOPS = I/O Operations per Second

SSD Concentrates Performance on Less Data

- Balance node CPU and Storage I/O with SSD impacts data space
  - > Same performance per Node
  - > Up to 4X more nodes for same data space as HDD only
  - > Up to 4X higher overall performance per TB
- Result: Flexible and efficient systems
  - > Match both performance and data space needs

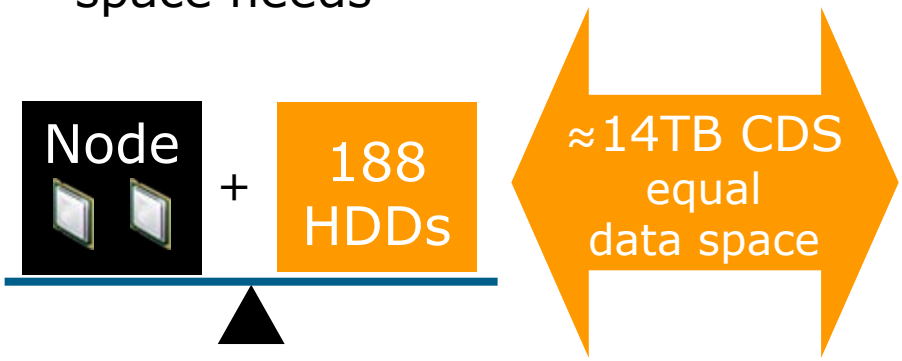

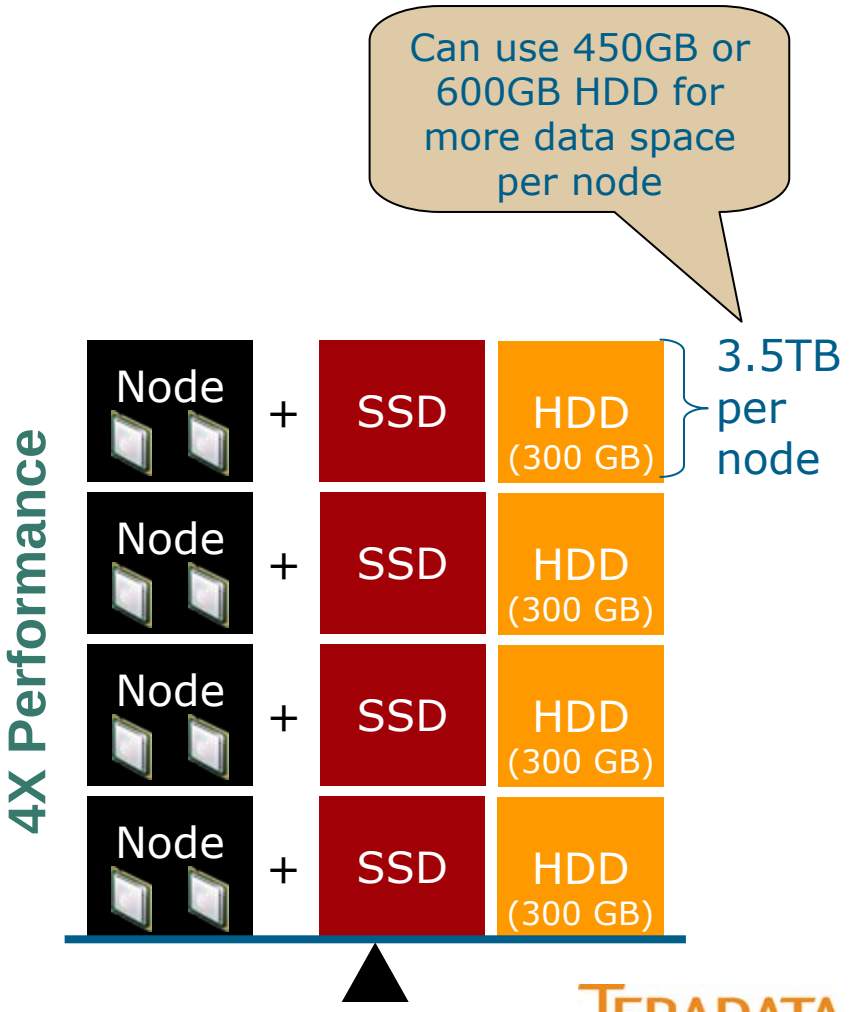

TERADATA
Raising Intelligence

Teradata Virtual Storage
Automatic Data Migration Across Drive Types

The next step in the Teradata Multi-Temperature Warehouse

- Drive Types are mixed within the system
  - > Fast solid state disks (SSD), high speed enterprise drives (HDD)
- Teradata Virtual Storage automatically places data on drive type that corresponds to the data temperature
  - > Optimum use of the storage resources – unique in the industry!

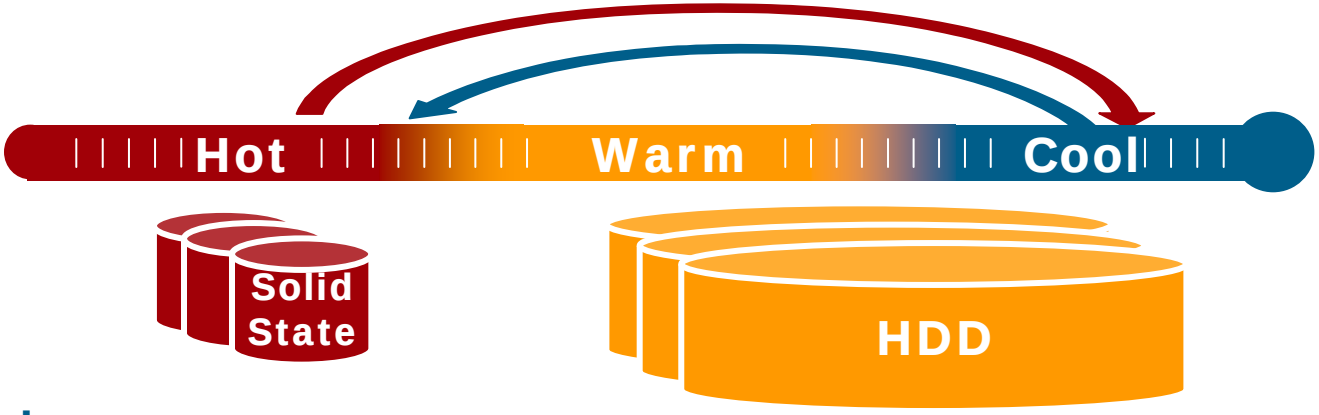

TERADATA
Raising Intelligence

Mixed Storage and Teradata Virtual Storage

- Optimizes hot data queries by migrating into high performance SSDs
- Cold data in HDD for lower performance reporting and history analytics
- Teradata optimizes use of large main memory in each node as Hot Cache

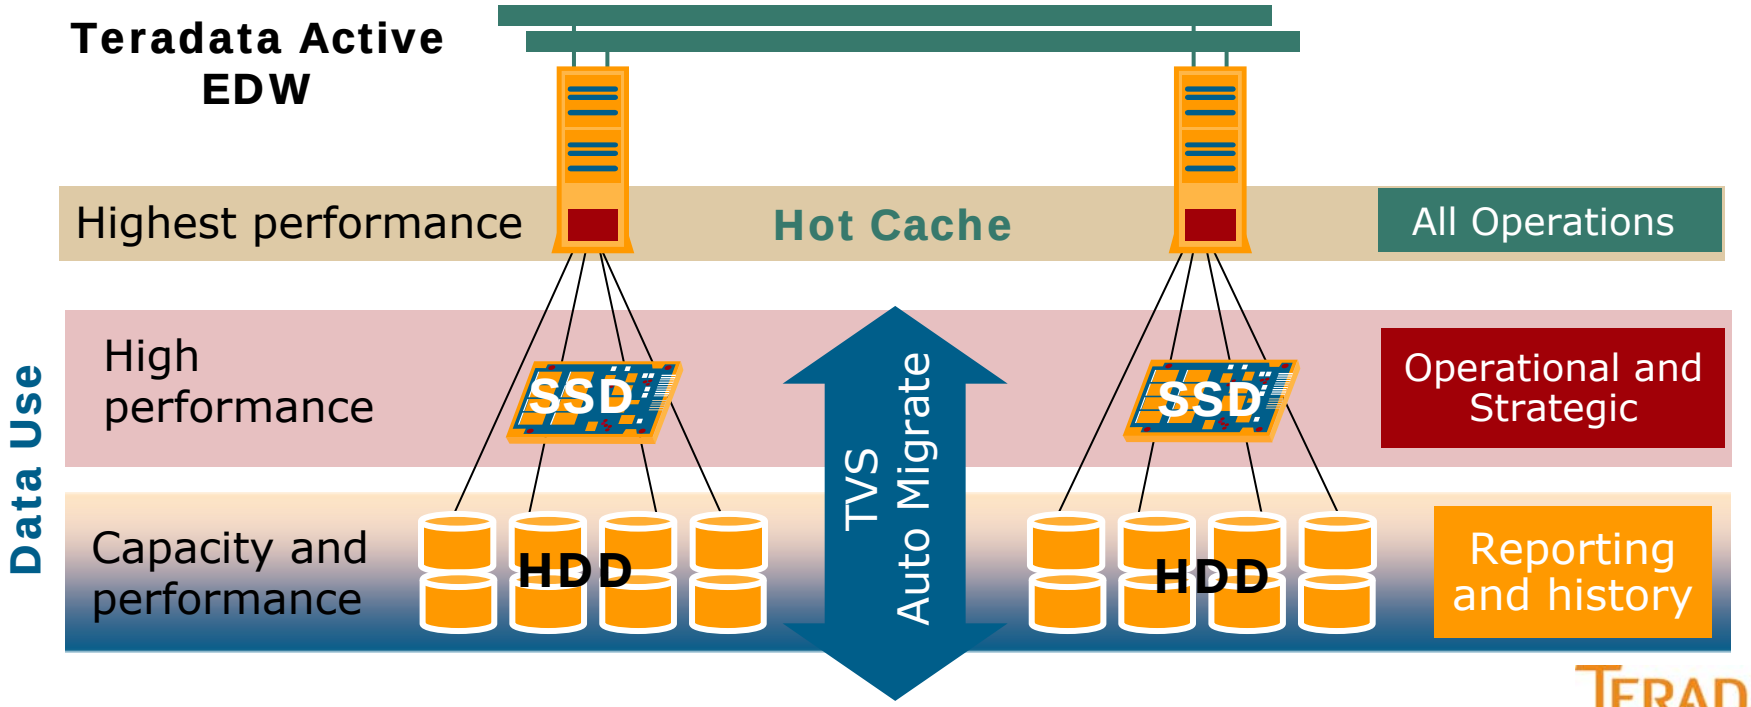

Temperature of Data is Set by Usage

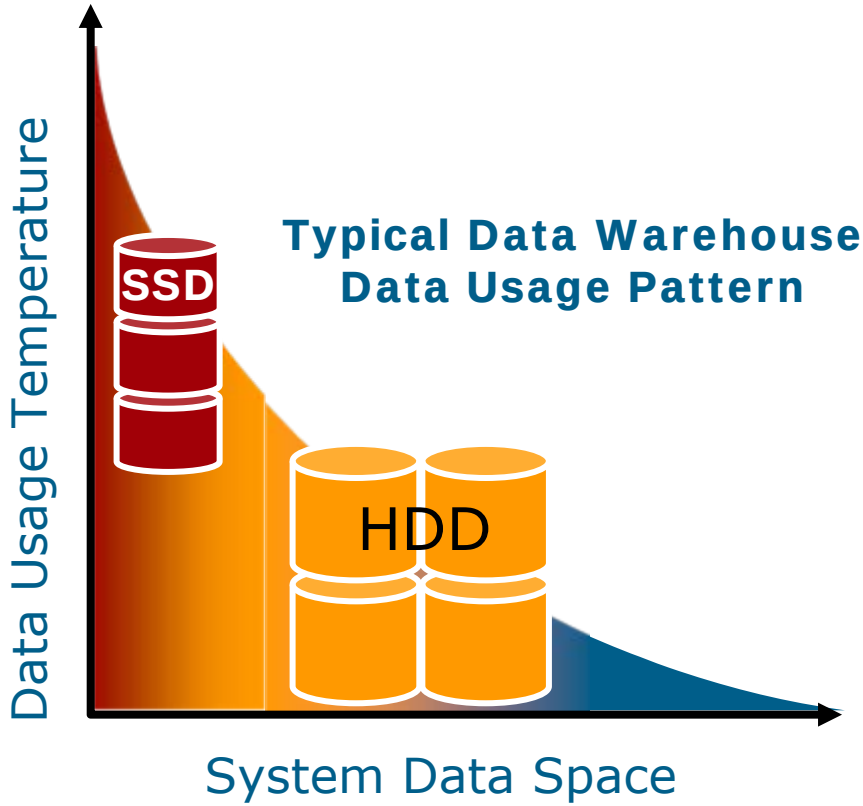

- ≈25% of EDW data is hot
  - > Used most frequently
  - > Recent data
  - > Last day, week, month

- ≈75% of data is warm/cold
  - > Accessed infrequently
  - > History – months ago
  - > Deep detailed info

Data Warehouse Workload Testing

6680: 100% of data on 12 SSD

vs.

5650: 100% of data on 188 HDD

Throughput Performance Improvements

Application	Test Query	1X 6680 node vs. 5650 node	4X 6680 nodes vs. 5650 node
Data Loading	MultiLoad/TPT average	1.15X	4.6X
	TPump/TPT stream average	1.12X	4.5X
Strategic W/L	Multi-Stream DSS average	1.05X	4.2X
Active Workload	Call center all AMP tactical	1.43X	5.7X

Data Space Equal

Faster & Consistent Query Response Times

- Average query times reduced for data on SSD
- SSD eliminates HDD delays
  - > Enables more consistent query times
  - > Narrows the response time curve

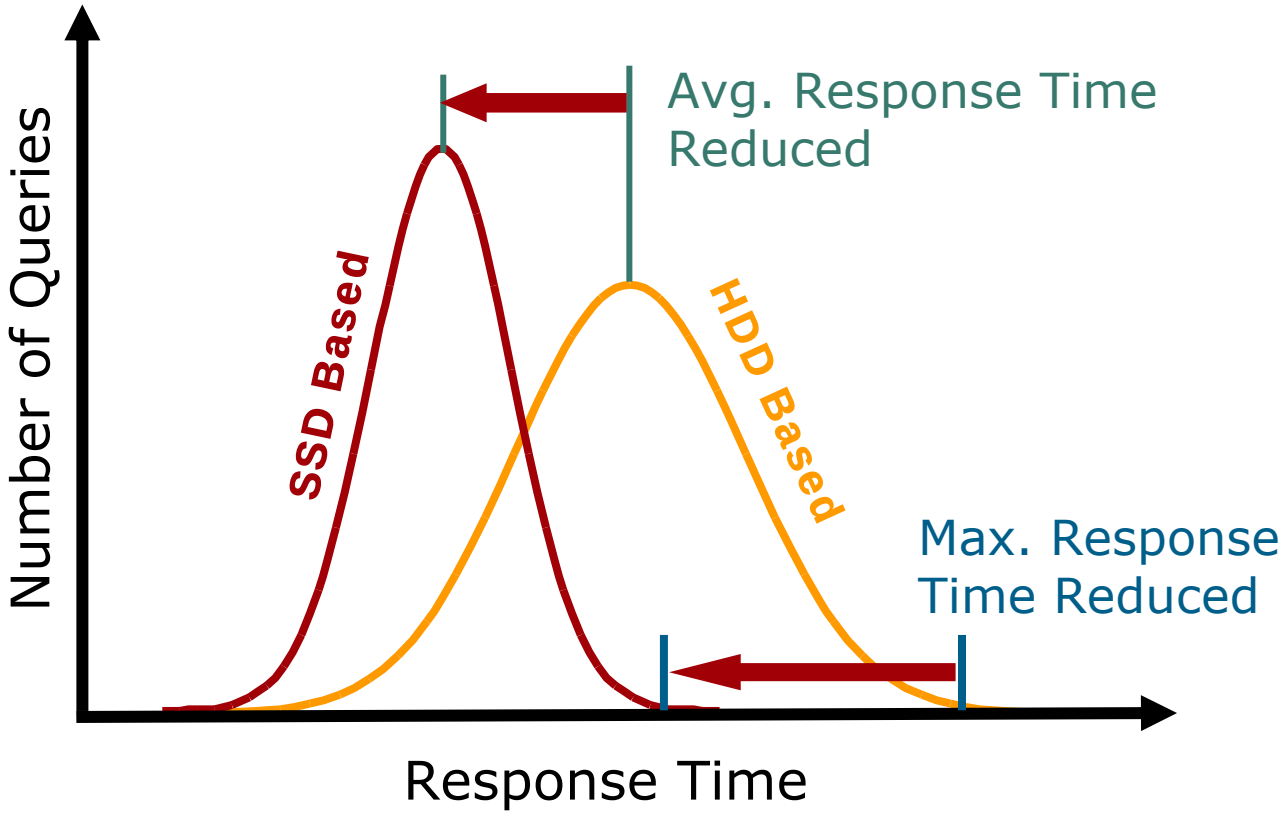

Teradata Active EDW 6650 & 6680 Solutions

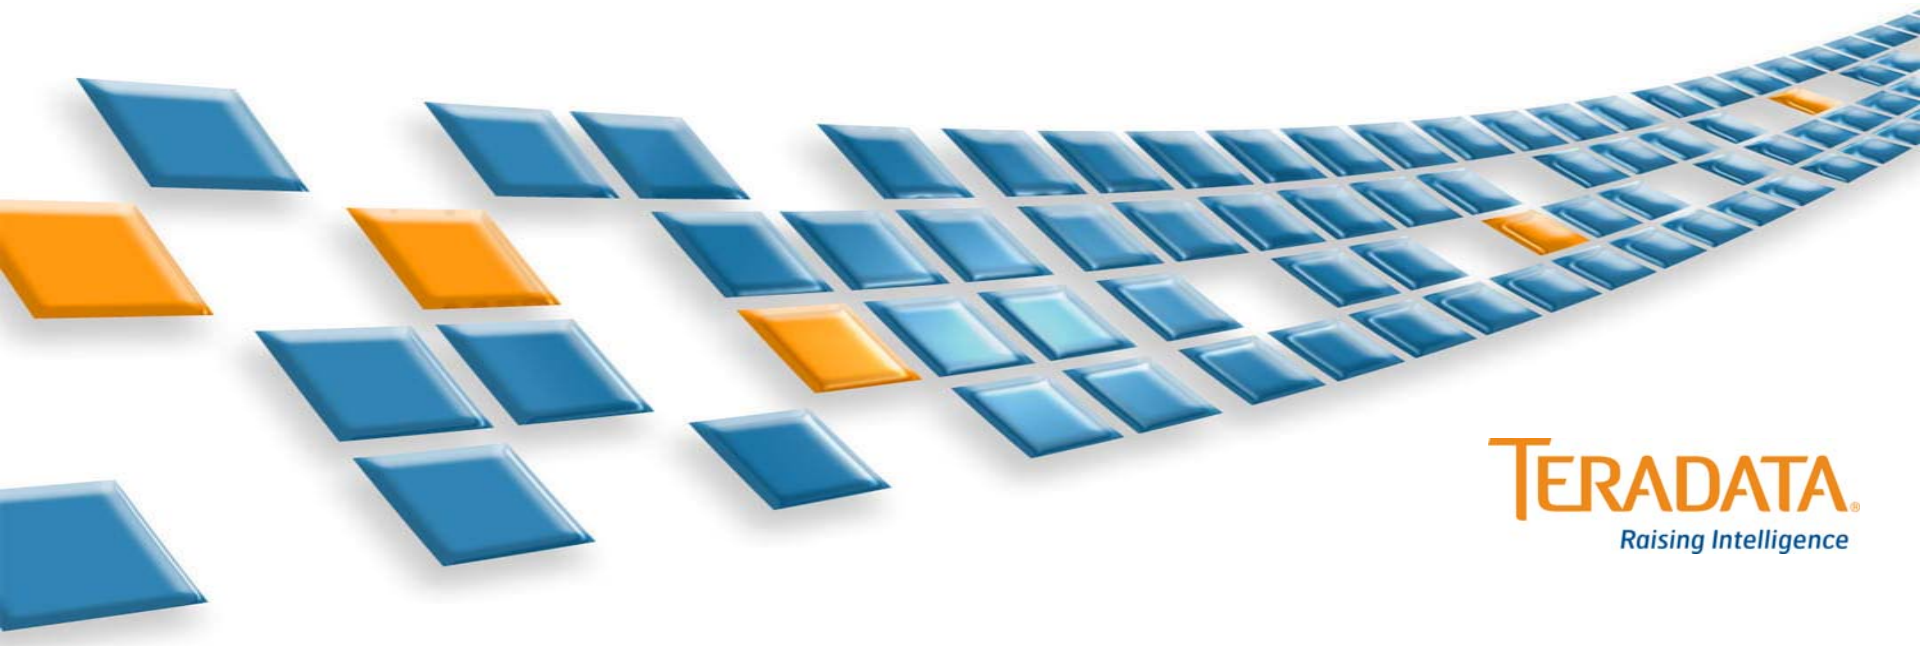

Teradata Active EDW: Use 6650 or 6680?

Growing performance at same pace as data storage

- Satisfied with performance
- Future high performance needs – add SSD later
- Coexistence with prior generations

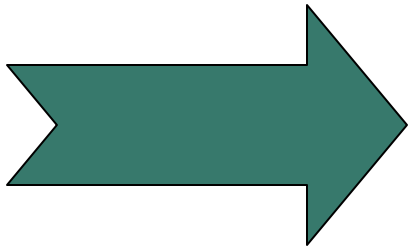

6650
HDD Only

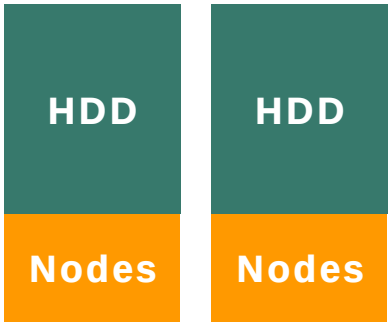

Growing performance at a higher pace than data storage

- Tighter SLA for Active workloads
- Floor space constrained
- Hot data sets with cold data
- Performance needs are NOW

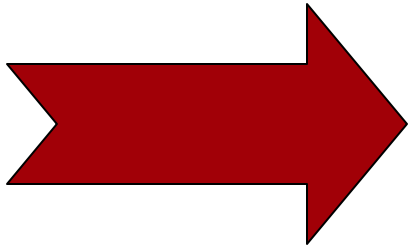

6680
Mixed Storage

TERADATA
Raising Intelligence

Teradata Active EDW 6650 – Matches the 5650
Saves Energy & Floor Space

5650 Clique

Basic building block

4 cabinets

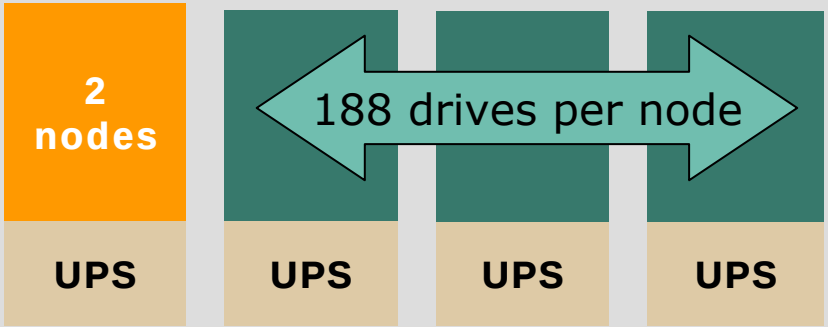

6650 Clique

Basic Building Block

3 cabinets

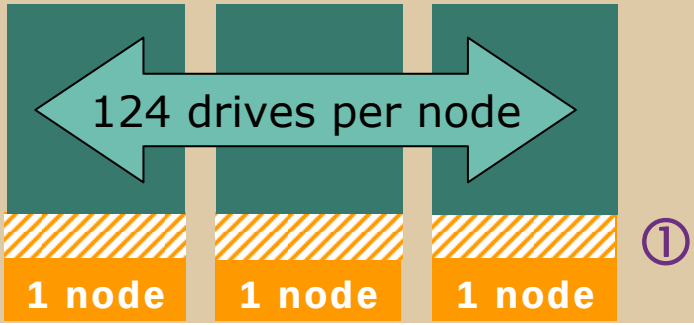

	5650	6650	Compare
User data space per clique	29.6TB	29.1TB	equal
HDD per clique	376	372	equal
Nodes per clique	2	3	
TPerf per clique	238	243	~equal
Energy & floor $/year	**$82K**	**$67K**	**-20%**

① **Node CPU capacity available for SSD add-in and special workloads**

TERADATA
Raising Intelligence

6680 Lowers Data Center Costs
Where Higher Performance per Data Space Needed

5650 System

Fill disks only 50% to get 2X performance per capacity

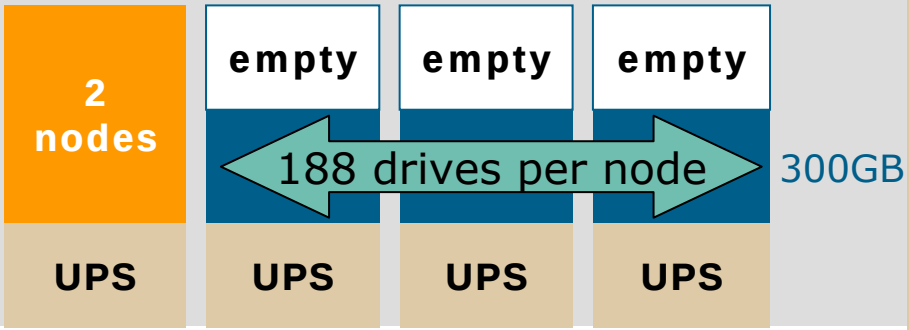

6680 System

SSD enables higher performance per data space

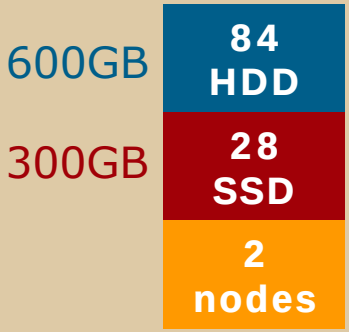

58% lower data center costs

	5650	6680	Compare
User Data Space delivered	29.6TB	15.6TB	Far fewer drives
User Data Space used	14.8TB	15.6TB	Save wasted space
Performance per data space*	**.96**	**1.06**	**1X**
Normalized list price	**$1**	**$.80**	**.8X**

*Normalized TPerf per TB

6680 - Efficient Approach to 2X Performance
with Same Data Space

6650 System

2 Nodes with Traditional Performance per Data Space

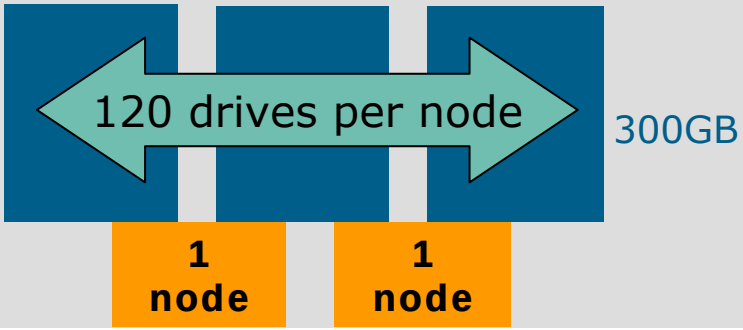

6680 System

4 Nodes with Same Data = 2X Performance per Data Space

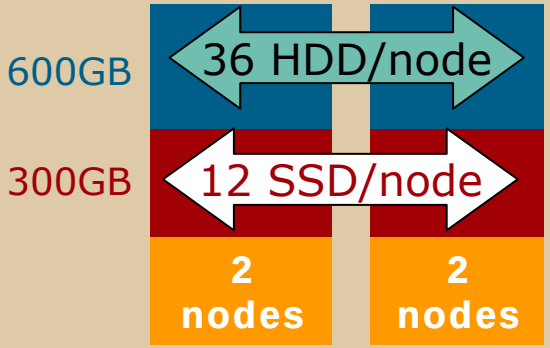

27% lower data center costs

	6650	6680	Compare
Data space	29.1TB	29.2TB	Same
TPerf per system	**243**	**504**	**2.1X**
Performance per Data Space*	**.49**	**1.04**	**2.1X**

*Normalized TPerf per TB

6680 - Efficient Approach to 4X Performance
with Same Data Space

5650 System

Traditional Performance per Data Space

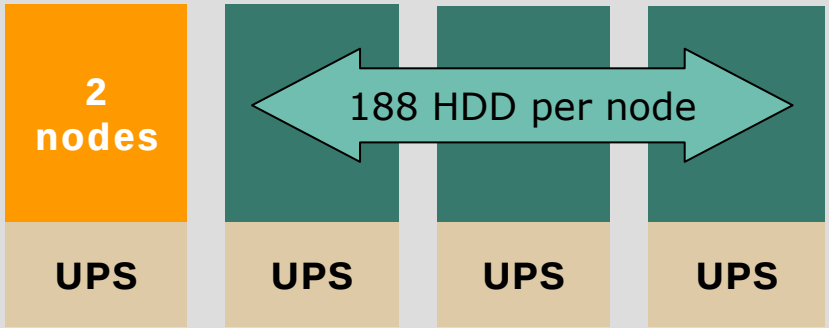

6680 System

4X Nodes @ Same Data Space =

4X Performance per Data Space

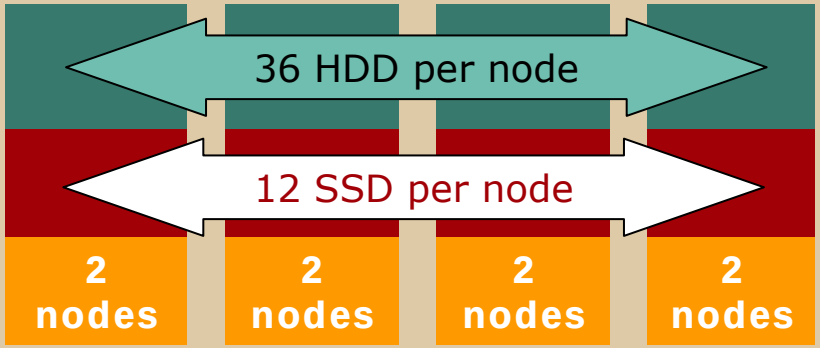

	5650	6680	Compare
Data space	29TB	30TB	**equal**
HDD/SSD capacity	300/NA	300/300	**same**
Performance per Data Space*	**.48**	**2.0**	**4X**

*Normalized TPerf per TB


TERADATA®
Raising Intelligence