

The image features a background of a globe with a network of glowing white lines connecting various points across the continents. A bright sun is positioned in the upper right quadrant, casting a strong glow and lens flare across the scene. The right side of the image has a blue background with a repeating pattern of small circles and vertical lines. In the top left corner, there is an orange rectangular box containing the Teradata logo.

TERADATA

# Teradata Database – Enabling Business Outcomes with Technology



**Customer**  
Experience



**Finance**  
Transformation



**Product**  
Innovation



**Risk**  
Mitigation



**Supply  
Chain**  
Intelligence



**Assets**  
Optimization



# Agenda

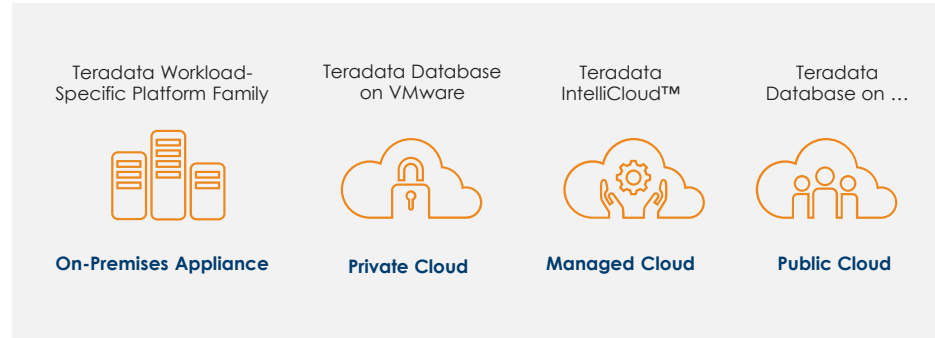
- New Bundling model for Teradata Database
- Shifting to new Cadence in our database release schedule
- New MAPS capability to provide higher performance and minimize downtimes








**Accelerate and Simplify**

# Teradata Everywhere

- Deployment flexibility
  - Fully featured Teradata Database
  - Same database across platforms
  - MPP scalability
- Design for data gravity
  - Co-locate with source data to minimize data movement and duplication
- Employ multi-system architectures
  - High availability, disaster recovery, workload distribution and shifting
- As your business changes, so can your deployment choice



# Teradata Everywhere drives new capability

	Cloud		On-Premises		
	Public	IntelliCloud	VMware	2xxx/ IntelliBase	6xxx/ IntelliFlex
					
Enterprise					
Advanced					
Base					Not Applicable
Developer		Not Applicable		Not Applicable	Not Applicable

# Bundled Features

Now develop how you want, deploy where you need

	Developer	Base	Advanced	Enterprise
Columnar	Included	Included	Included	Included
Temporal	Included	Included	Included	Included
Row-level Security	Included	Included	Included	Included
Secure Zones	Included	Included	Included	Included
Workload Management	Not Applicable	Not Applicable	TIWM	TASM

## If the business needs:

- Rapid performance on selective columns
- If your business captures lots of attributes but uses selectively or in groups for specific applications
- If data is fairly static with targeted access
- Performance relief for large user communities or I/O constrained systems
- Chargeback relief of data storage with automatic compression



**Teradata Columnar**

# Columnar Examples

- Telecom
  - Call Detail Records stored in wide fact table with hundreds of columns
  - Analytic queries typically hit a subset of columns, 10 -12
- Airline / Transportation
  - PNR has extensive detailed for each passenger
  - Often only need Date, Origin, Destination
- Health Care
  - EHR contains details about vitals and visit care
  - Typically need patient and code for analytics
- Results:
  - Significantly reduces I/O and saves CPU cycles
  - Dramatically speeds response time
  - Reduces storage needs with auto-compression
  - Gets even better results if combined with row partition elimination
  - Allows data to be kept integrated without performance degradation



# Columnar Basics

- This is an optimization technique, of how the database manages the data. It is not a “query feature”
- Hybrid Columnar, not an all or nothing approach
- Can be used in conjunction with Row Partitioning
- Can create JI with Columnar format on a row base table
- Major enhancement in 15.10
  - NoPI restriction lifted
  - Qualified update or delete in place
- Like any other optimization technique, there are pros and cons
- Implemented in the “Create Table” statement

## If the business needs:



- Ability to compare past, present and future patterns of changing hierarchies
- Ability to align multiple hierarchies and data changes to run analytics
- Ability to understand and analyze change over time
- Ability to certify and prove what was known and when it was known
- Ability to go beyond just “begin and end dates” but also understand effective dates to changes






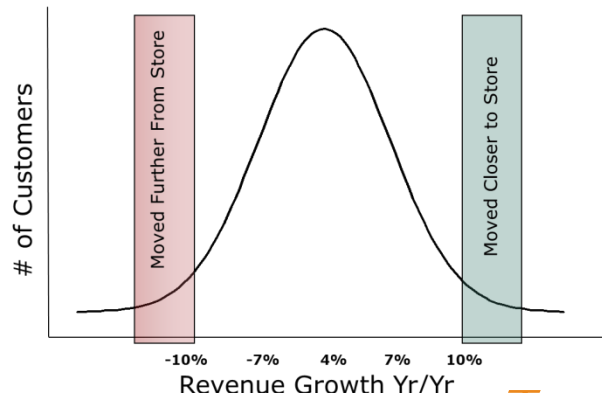
**Teradata Temporal**

# Temporal Examples

<p><b>July 1, 2008</b></p>  <ul style="list-style-type: none"> <li>• New hire Frank starts work</li> <li>• Chooses premium insurance policy</li> </ul>	<p><b>September 18, 2009</b></p>  <ul style="list-style-type: none"> <li>• Frank gets hurt on job</li> <li>• Starts rehab program</li> </ul>
<p><b>January 1, 2010</b></p>  <ul style="list-style-type: none"> <li>• Frank back on job</li> <li>• Changes to basic insurance policy for 2010 at open enrollment to save money</li> </ul>	<p><b>February 12, 2010</b></p>  <ul style="list-style-type: none"> <li>• Injury rehab claim submitted</li> <li>• Frank currently shows basic insurance policy</li> <li>• Which policy terms are used in claims processing?</li> </ul>

<p><b>January 1, 2010</b> Sales Regions</p> 	<p><b>February 28, 2010</b> YTD Sales</p> <table border="1"> <thead> <tr> <th></th> <th>YTD Sales</th> <th>2010 Q1 Goal</th> </tr> </thead> <tbody> <tr> <td>West</td> <td>\$124M</td> <td>\$160M</td> </tr> <tr> <td>East</td> <td>\$180M</td> <td>\$ 215M</td> </tr> </tbody> </table>		YTD Sales	2010 Q1 Goal	West	\$124M	\$160M	East	\$180M	\$ 215M
	YTD Sales	2010 Q1 Goal								
West	\$124M	\$160M								
East	\$180M	\$ 215M								
<p><b>March 1, 2010</b> Sales Regions</p> 	<p><b>March 31, 2010</b> YTD Sales</p> <table border="1"> <thead> <tr> <th></th> <th>YTD Sales</th> </tr> </thead> <tbody> <tr> <td>West</td> <td>\$215M</td> </tr> <tr> <td>East</td> <td>\$175M</td> </tr> </tbody> </table>		YTD Sales	West	\$215M	East	\$175M			
	YTD Sales									
West	\$215M									
East	\$175M									

<p><b>2009 Sales</b></p> <table border="1"> <thead> <tr> <th></th> <th>2009 Sales</th> <th>2010 Goal for Bonus</th> </tr> </thead> <tbody> <tr> <td>Frozen Foods</td> <td>\$124M</td> <td>\$136M</td> </tr> <tr> <td>Dairy</td> <td>\$ 75M</td> <td>\$ 90M</td> </tr> </tbody> </table>		2009 Sales	2010 Goal for Bonus	Frozen Foods	\$124M	\$136M	Dairy	\$ 75M	\$ 90M	<p><b>2010 Sales</b></p> <table border="1"> <thead> <tr> <th></th> <th>2010 Sales</th> </tr> </thead> <tbody> <tr> <td>Frozen Foods</td> <td>\$126M</td> </tr> <tr> <td>Dairy</td> <td>\$ 92M</td> </tr> </tbody> </table> <p><b>BONUS</b></p>		2010 Sales	Frozen Foods	\$126M	Dairy	\$ 92M
	2009 Sales	2010 Goal for Bonus														
Frozen Foods	\$124M	\$136M														
Dairy	\$ 75M	\$ 90M														
	2010 Sales															
Frozen Foods	\$126M															
Dairy	\$ 92M															
<p><b>February 13, 2010</b> Item Category Change</p> <table border="1"> <thead> <tr> <th>Frozen Foods</th> <th>Dairy</th> </tr> </thead> <tbody> <tr> <td></td> <td></td> </tr> <tr> <td></td> <td>Ice Cream</td> </tr> </tbody> </table> <p>2010 Ice Cream Sales \$12M</p>	Frozen Foods	Dairy				Ice Cream	<p><b>2010 Sales</b> (Categories as of Dec. 31, 2009)</p> <table border="1"> <thead> <tr> <th></th> <th>2010 Sales</th> </tr> </thead> <tbody> <tr> <td>Frozen Foods</td> <td>\$138M</td> </tr> <tr> <td>Dairy</td> <td>\$ 80M</td> </tr> </tbody> </table> <p><b>BONUS</b></p>		2010 Sales	Frozen Foods	\$138M	Dairy	\$ 80M			
Frozen Foods	Dairy															
																
	Ice Cream															
	2010 Sales															
Frozen Foods	\$138M															
Dairy	\$ 80M															




# Temporal Basics

- Automatically manages time periods with table updates
- Simplifies ETL coding process
- Provides Bi-temporal support
- Simplifies queries for time based analytics
- “Normalizes” adjacent ranges
- Improves quality and consistency of data and queries
- Retains historical context for business operations
- Talk to your BI vendor for further support

## If the business needs:

- Distinct data visibility based on user certification
- Requires Mandatory Access Control (MAC) compliance
- Security at the data row level not via views
- Additional auditability



**Row Level Security**



# Row Level Security Examples

- Security Labels

- Associated with subjects (users) and objects (tables, rows)

- Two parts

- *Classification* - a single, hierarchical level

- e.g., military classification: TOP SECRET, SECRET, CONFIDENTIAL, UNCLASSIFIED

- *Compartments* – (optional) nonhierarchical - represent distinct areas of information


- e.g., Region or Department: HR, FIN, SALES, MKT, CS

**Label**

**Users**

**Rows (Data)**

**Label**


SECRET [HR]			SECRET [HR]
			CONFIDENTIAL [HR]
			CONFIDENTIAL [HR, FIN]
CONFIDENTIAL [FIN]			SECRET [FIN]
			TOP SECRET [HR]

# Row Level Security Basics

- Need to consider View usage versus row level usage
- Requires UDF to be created which defines handling rules
- Tables need new column to contain labels
- CAUTION: Results may vary depending on what you can see
- Better logging, reporting, and auditing than with views

## If the business needs:

- Security between user groups and data content
- Isolation from administrative role versus access and analysis role
- To host Multiple independent applications within single platform
- “Virtual Data Warehouses” with controlled cross data access
- Physically, not logically, manage data tables

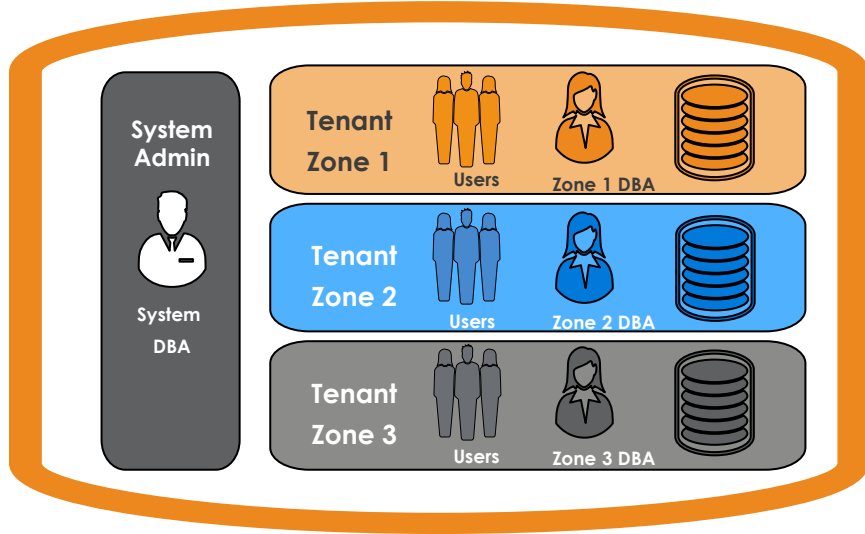


**Secure Zones**

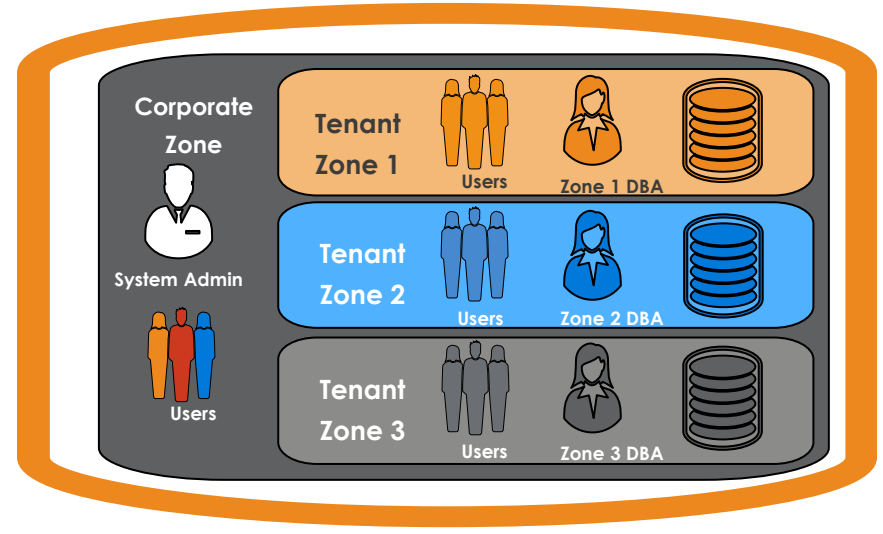
# Secure Zone Examples



## Multi Tenant Hosting



## Corporate Isolation



# Secure Zone Basics

- Global uniqueness on naming objects (this is being addressed)
- Users are assigned to single zone but can have “guest access” to other zones
- About access not resources
- Integrated with Studio for administration tasks
- Combine with TASM / TIWM for complete management



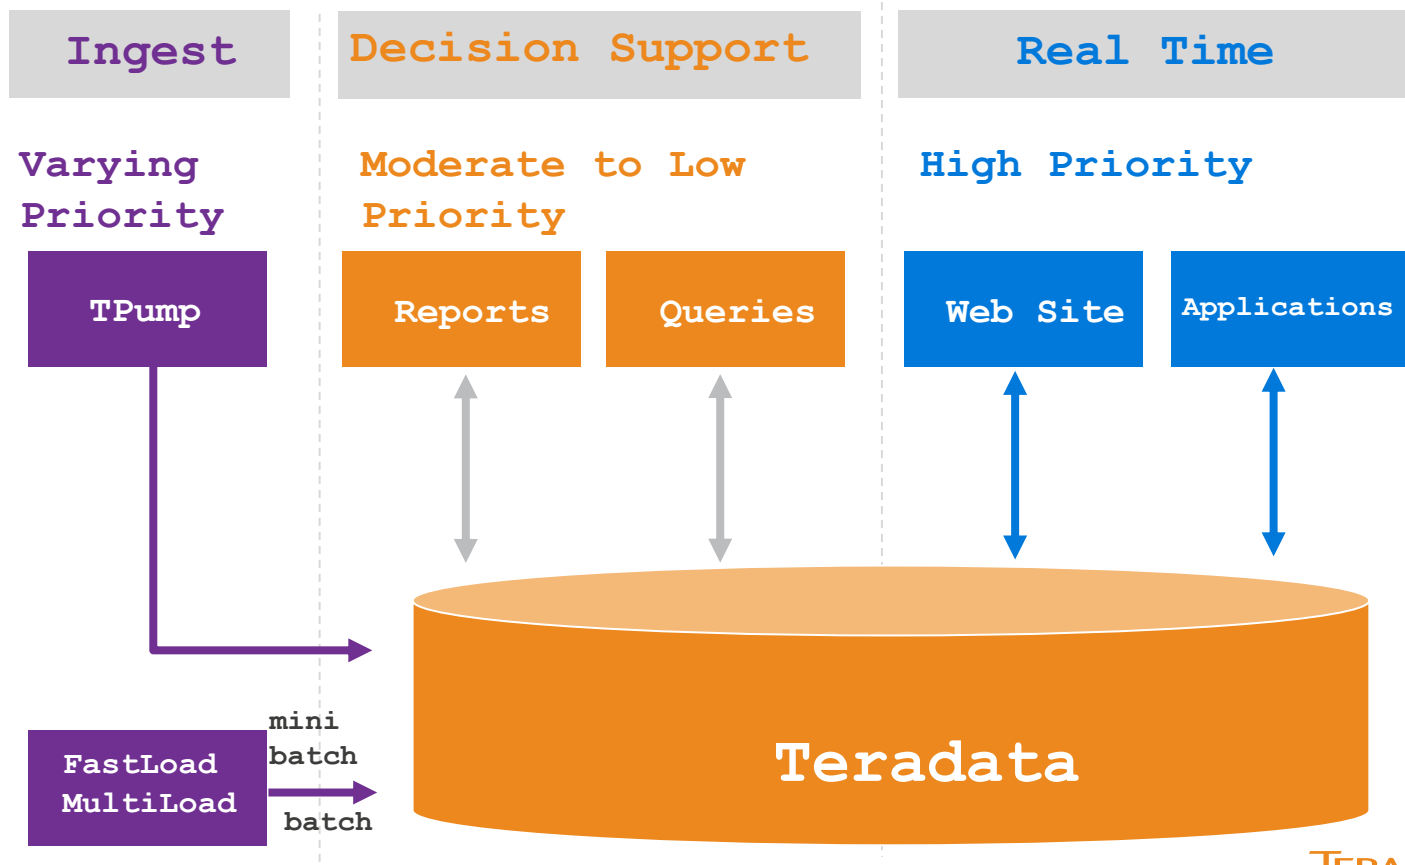
## If the business needs:

- Mixed workloads with consistent response times
- Tactical workloads with guaranteed resources available
- Virtual platforms with defined resources commitments
- Periodic or Scheduled prioritization



**Workload Management**

# Teradata Workload Management Example



# Workload Management Basics

Teradata Workload Management Features	Teradata Integrated Workload Management	Teradata Active System Management (TASM)
Workload Classification	Source, Target, Query Characteristics, QueryBand, Utility	
Workload Priority Management	Tactical and Timeshare	Tactical, <b>SLG Tiers</b> , and Timeshare
Resource Management	CPU and I/O and Memory	
Virtual Partitions	One Virtual Partition	<b>Multiple</b> Virtual Partitions
Filters and Throttles	Filters, Workload and System Throttles	<b>Flex Throttles</b>
Exceptions	Tactical Exceptions and Timeshare Decay	Tactical and <b>Workload Exceptions</b> , and Timeshare Decay
State / State Matrix	Operating Periods	Operating Periods and <b>Health Conditions</b>

# Getting Started --

- Be aware of business needs
  - Fit into current environment
  - Compliment with new feature
- Prove the value
  - Define the “quick” and the “win”
  - Key user group with real pain
- Document and Market
  - Highlight the success
  - Measure productivity, cost, or time improvements
- Repeat and Extend
  - Add a new level of complexity
  - Extend to new users and use cases




# Further Reading and References

- Teradata.com
  - [William McKnight White Paper](#)
  - [Video: A Smarter Way To Do Analytics](#)
  - [Teradata.com page on columnar](#)
  - [Enabling Temporal White Paper](#)
  - [Article: Multiple Data Warehouses in One System](#)
  - [Smart Recommendations for Small Businesses at Network Solutions®](#)
  - [On the fast track towards advanced analytics](#)
- Techbytes
  - [Secure zones](#)
- DevX or Partner
  - [Temporal and OBIEE](#)
  - [Temporal and Microstrategy](#)
- TEN (Teradata Education Network)
  - [WBT: Introduction to Temporal](#)
  - [Webcast: Advanced Temporal](#)
  - [ILT: History, Time Variance, and Temporal](#)
  - [Teradata Secure Zones: Implementation Basics](#)
  - [Workload Management Tips and Techniques](#)
  - [Workload Management Technical Overview](#)
- T@YS - Orange books
  - 541-0009036-B02 TERADATA COLUMNAR FOR TERADATA DATABASE 15.10
  - 541-0010639A02 TERADATA SECURE ZONES
  - 541-0009217-A03 TERADATA ROW LEVEL SECURITY
  - TDN0001776 TERADATA ACTIVE SYSTEM MANAGEMENT TERADATA DATABASE 16.0
  - TDN0001728 - 3.0 TERADATA INTEGRATED WORKLOAD MANAGEMENT TERADATA DATABASE 16.0



## Giving the Business:


- ~ 90% reduction in system expansion downtime
- Cloud elasticity; expand/shrink your Teradata system
- More consistent tactical query performance
- Higher query concurrency



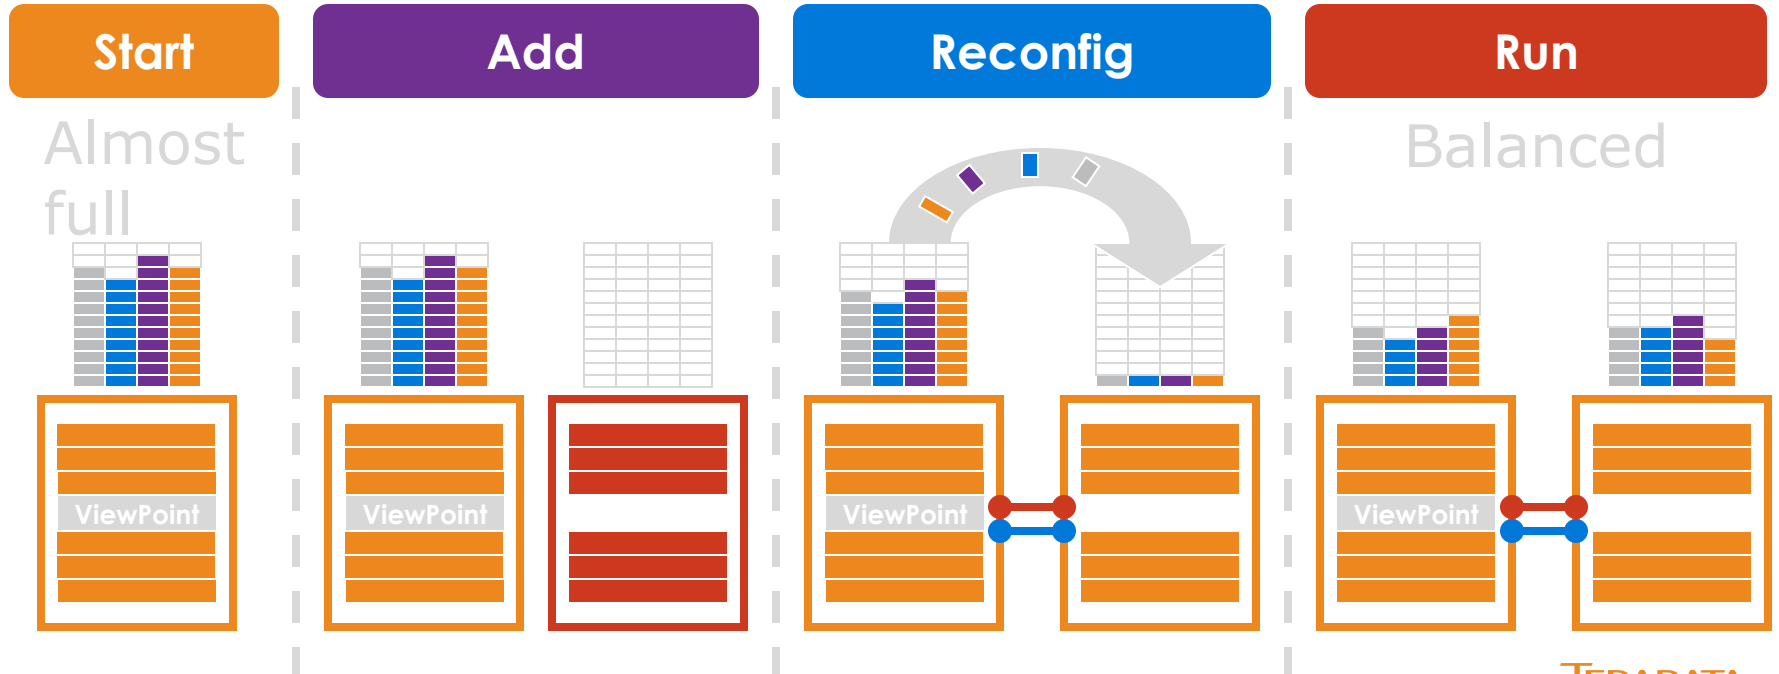
**Teradata Maps**

# Hash Map

- Maps (i.e., translates) which hash buckets go on each AMP

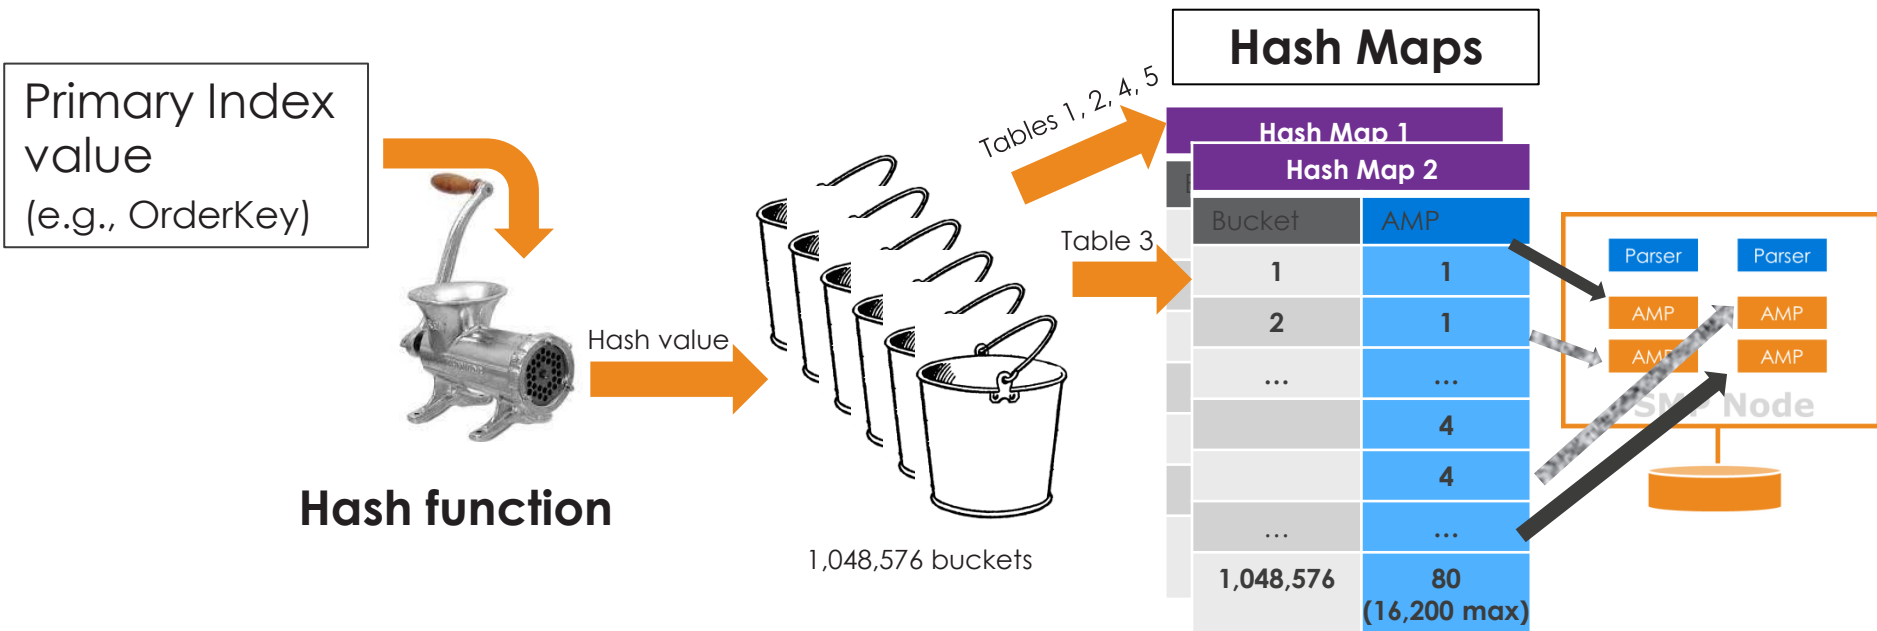


# Expanding the System



# MAPS: Enabling Multiple Hash Maps

- Maps (i.e., translates) which hash buckets go on each AMP
- Use different Hash Maps for different tables or table partitions






# MAPS Provides New Flexibility

- Very large system expansion reconfigs
  - High availability and reduced planned outage by scheduling movement of data to expansion hardware
  - Move only specified data when desired
  - Complete flexibility on how data is spread for optimum performance, system utilization, and availability for business needs
- Small tables (number of rows relative to number of AMPs)
  - Inefficient to involve every AMP for small table or few rows/AMP



# Sparse Map



TD\_Map1 – Contiguous Map



Table X → TD\_Map1




My\_5AMP\_Map – Sparse Map



Table Stores → FiveAmpMap

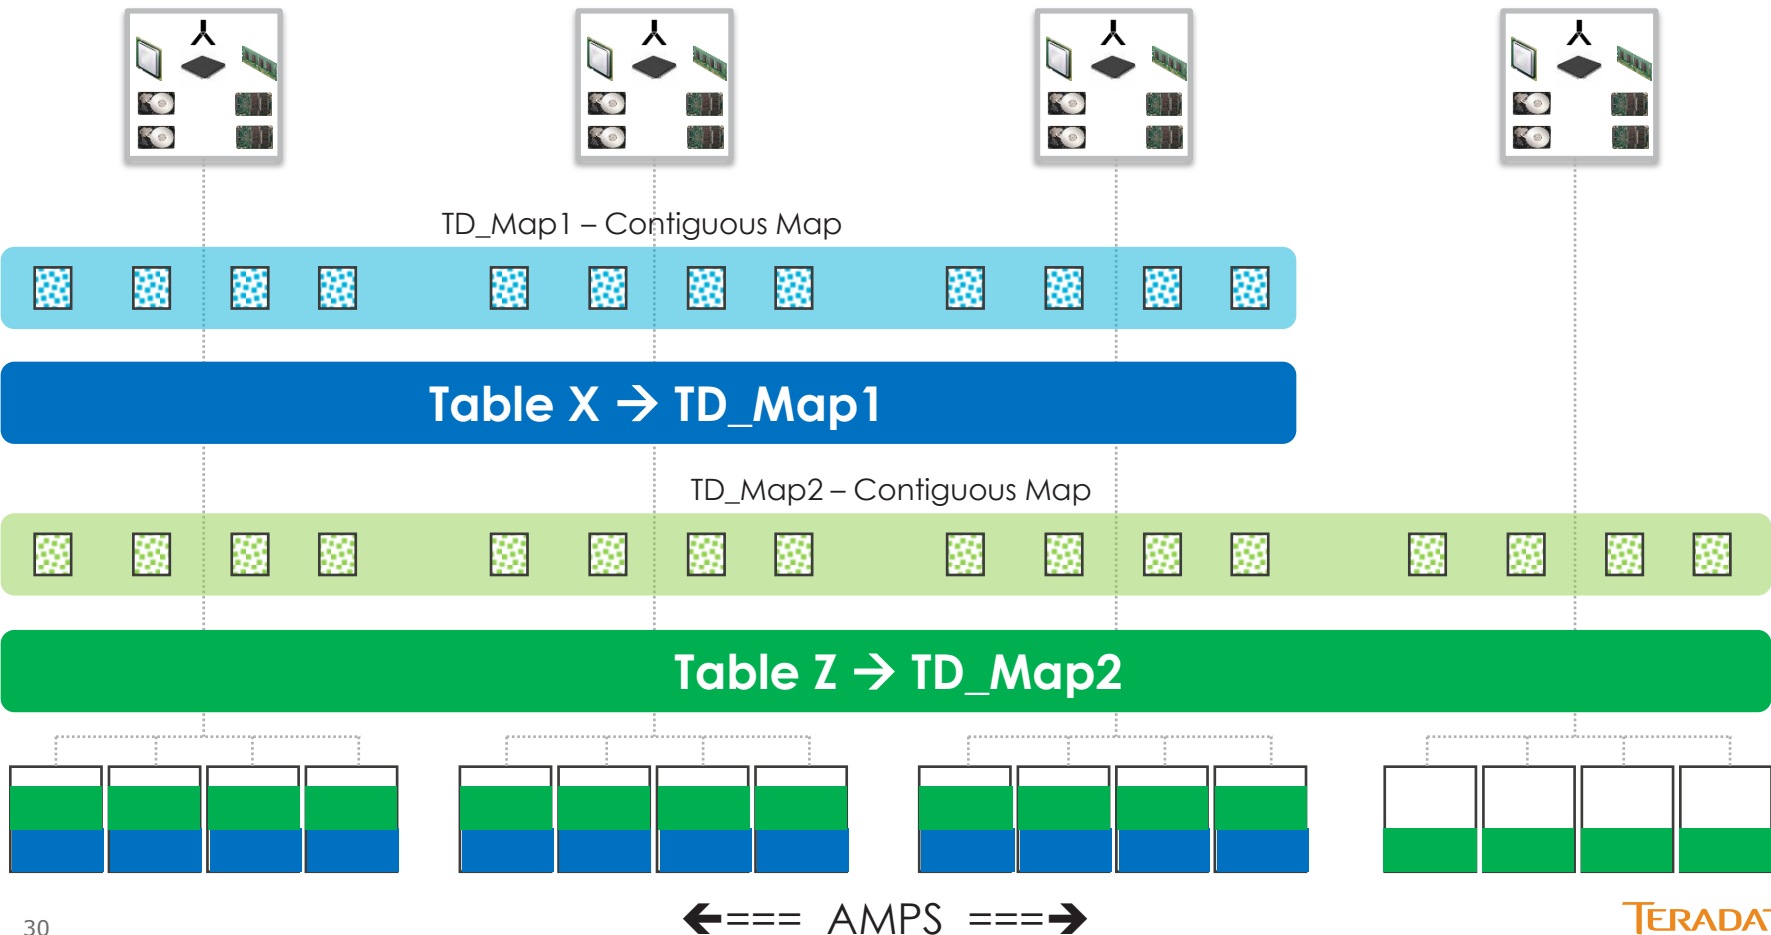
```
CREATE TABLE Stores,  
MAP = FiveAmpMap as StoreMap  
(column definitions)  
PRIMARY INDEX (Store_Id);
```

- Table co-location
- ✓ same sparse map
  - ✓ same map alias\_name
  - ✓ same Primary Index



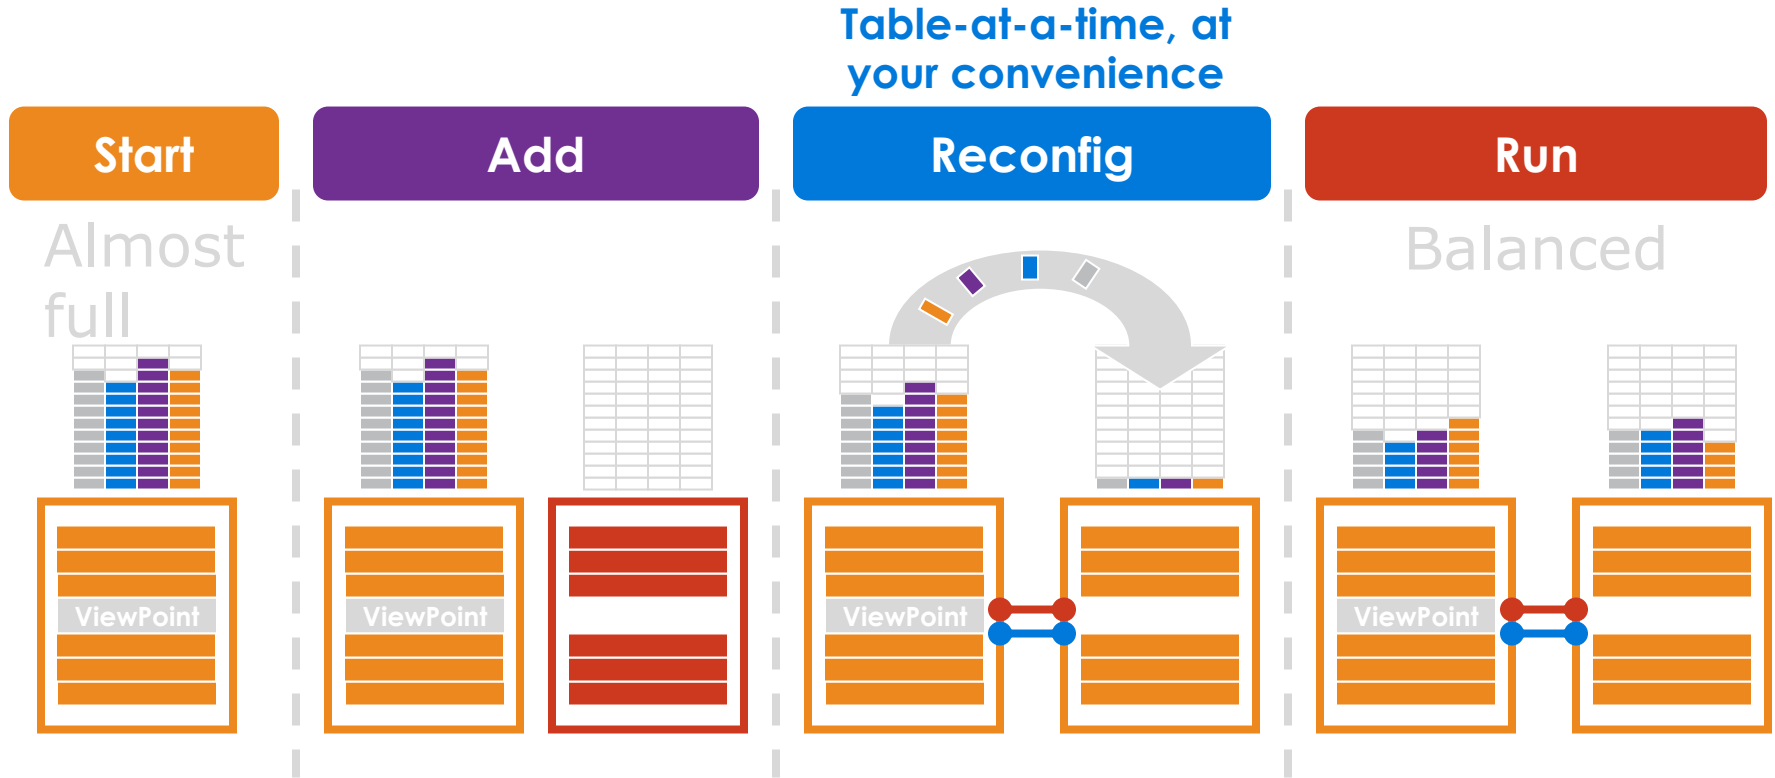
←=== AMPS ===→

# Expansion



# MAPS Provides System Flexibility and Agility

## *Expanding the System*



# Giving the Business:

- New features available faster
  - No longer have to wait for major feature and 12-18 month release cycle
  - Customer feature enhancement requests can be accommodated sooner
  - Database capabilities stay leading edge delivered in smaller increments
- Upgrade planning becomes more predictable and stable
  - More frequent and consistent release schedule
- Reduced risk when adopting a critical feature
  - Smaller amount of change in EACH release reduces risk and makes testing easier
- At forefront of industry direction



**Teradata Release  
Cadence**

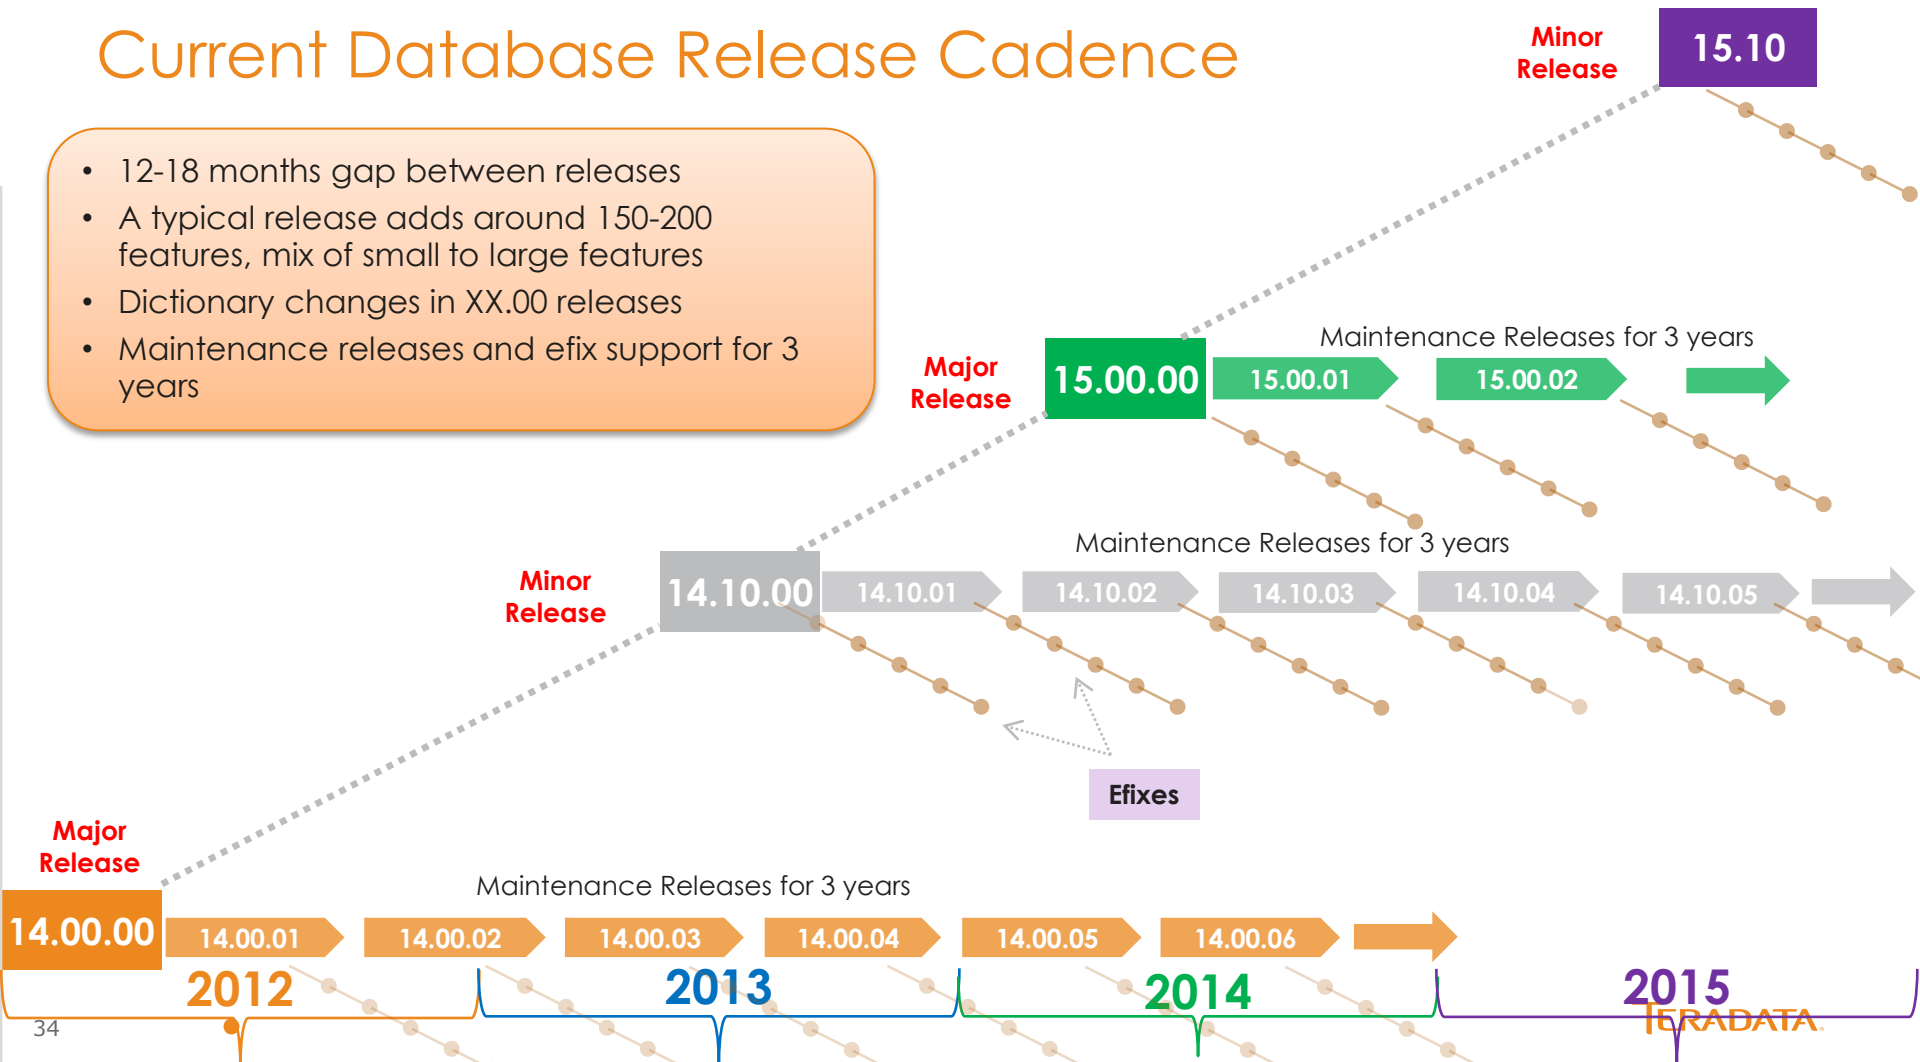
# Teradata Labs and Teradata Database Becoming More Agile

- Teradata Database, Teradata Tools and Utilities, and most UDA Ecosystem releases generally every four months beginning in 2017
- Product planning and requirements process becoming agile




# Current Database Release Cadence

- 12-18 months gap between releases
- A typical release adds around 150-200 features, mix of small to large features
- Dictionary changes in XX.00 releases
- Maintenance releases and efix support for 3 years




# Agile Evolution

## Migrating to a Four Month Release Cadence





# Four Month Database Release Cadence



# Teradata Database Implementation Details

- Frequent releases
  - Generally every 4 months
  - Some cycle variation
- All releases have same “importance” (e.g., second as significant as others)
- XX.00 release at beginning of every other release series has “major release restrictions”
  - Possible data dictionary changes
  - No back-down once upgrade to this release
- Efix/Maintenance releases
  - XX.00, XX.10, XX.50, XX.60 releases (first two releases of series) receive efixes for 4 months (or until next release GCA if longer)
  - XX.20 and XX.70 releases (last release of series) receive maintenance releases and efixes for 3 years
  - Updated support policy posted
    - refer to the Teradata Support Policy “Teradata Database Software Support Lifecycle” section <http://www.teradata.com/Teradata-Product-Support-Policies/>
- Simplifying Beta process
  - Customer can participate in Beta for future releases without repeat contract work

TERADATA®

# Release Cadence Back-up Detail

# Teradata Tools and Utilities (TTU) Implementation Details

- Released in concert with each Teradata Database release
  - Generally every four months
  - Monthly rollup maintenance releases
- Release numbers consistent with Teradata Database releases
  - In exception cases of additional minor TTU release, second point digit used (e.g., 16.11)
- Forward and backward compatibility with Teradata Database
  - No change from previous policy
  - Forward and backward compatible with whole Teradata Database release series
  - Partner certification with the latest TTU release will suffice for the series

# TTU Compatibility With Teradata Database

- 2 forward and 4 backward major or minor Teradata Database release series
  - Consistent with current policy
- TTU 15.10
  - TD13.10, 14.00, 14.10, 15.00, **15.10**, 16.00-16.20, 16.50-16.70
- TTU 16.00/16.10/16.20
  - TD14.00, 14.10, 15.00, 15.10, **16.xx**, 16.50-16.70, 17.00-17.20
- TTU 16.50/16.60/16.70
  - TD14.10, 15.00, 15.10, 16.00-16.20, **16.xx**, 17.00-17.20, 17.50-17.70
- TTU 17.00/17.10/17.20
  - TD15.00, 15.10, 16.10-16.20, 16.50-16.70, **17.XX**, 17.50-17.70, 18.00-18.20

# UDA Ecosystem Product Implementation Details

- UDA Ecosystem Products
  - Data Mover, Ecosystem Manager, Studio, Viewpoint, Data Lab, Data Stream Architecture (DSE/DSU), Unity
- Releases with each Teradata Database release
  - Generally every four months
  - Maintenance releases and efixes as necessary
  - QueryGrid is distinct
    - Remains on separate quarterly release schedule
    - Remains on separate release number sequence
- Release numbers consistent with Teradata Database releases
- Maintenance Policy
  - Policy remains the same
  - *Maintenance releases and efixes for these Products are only supplied on the most current UDA product released version. Customers on earlier versions should upgrade to the most current application version to obtain fixes.*



# UDA Products Compatibility With Teradata Database

- UDA compatibility with Teradata Database remains the same
  - “Current database + 4-back”
  - “Current” however is now defined as a series (e.g., 16.00, 16.10, 16.20)
- UDA Product 16.00
  - TD 14.00, 14.10, 15.00, 15.10, 16.00
- UDA Product 16.10
  - TD 14.00, 14.10, 15.00, 15.10, (16.00, 16.10)
- UDA Product 16.20
  - TD 14.00 14.10, 15.00, 15.10, (16.00, 16.10, 16.20)
  
- UDA Product 16.50
  - TD 14.10, 15.00, 15.10, 16.20, 16.50
- UDA Product 16.60
  - TD 14.10, 15.00, 15.10, 16.20, (16.50, 16.60)
- And so on ...

# Partner Validation of New Teradata Release Cadence

- ISV partners expected to perform validation testing on at least one Teradata Database release per year
- ISV partners expected to declare interoperability support for all releases in the cadence series (e.g. 16.00-16.20, 16.50-16.70, 17.0-17.20)
  - Enabled by two Teradata policies
    - Client software to be compatible across all database releases in a series
    - Teradata Database to maintain SQL and data dictionary compatibility throughout a release series
- Partners to support all Teradata Database deployment options (e.g., appliance, cloud)
  - Enabler: Teradata Everywhere
- Teradata expects ISV partners will increase validation testing frequency
  - Long-term goal
  - Automation and simplification of test processes