



Customer Best Practices for Teradata Upgrades

Major Hardware Upgrades

Revision History

Summary of Changes			
Date	Description of Change	Rev #	Author
August 15, 2008	Version for publication	1.0	SFT Members

Important Notes	
Date	

Table of Contents

1.	Introduction	1
2.	The Upgrade Process for Teradata Hardware.....	3
	Upgrade Process Overview Diagram	3
3.	Project Management	4
4.	Upgrade Planning	4
	Plan for capacity.....	4
	Research environmental requirements	4
	Review network requirements.....	5
	Pre-upgrade Benchmarking	5
5.	User Communication.....	6
6.	Change Control	6
7.	Environment Changes	7
	Data Center:	7
	MVS Environment:.....	7
	Node OS Environment:	7
	Network Environment:.....	7
8.	BAR Implications	8
9.	Pre-Upgrade Preparations with Teradata	8
10.	Teradata Upgrade	10
11.	Post-Upgrade Steps	11
12.	Project Closeout	12
13.	Closing	12

1. Introduction

A successful upgrade of any computing environment requires careful planning, coordination, and execution. A Teradata computing environment is no different. There are many factors that need to be considered and brought together for a successful plan. The success you experience is in part dependent on the success of your efforts throughout the process.

In general, there are three things that influence the success or failure of an upgrade:

- The quality of the product being added/updated
- The service procedures followed by Teradata associates
- The planning and execution of operational procedures by the customer

All three need to come together without flaw during the upgrade process in order to achieve success. The first two are the responsibility of Teradata. For the purposes of this paper, they are considered out of scope, but assumed to be present and bear no negative impact on the upgrade process.

In the case of hardware upgrades and expansions, it is also critical that the customer provide Teradata with enough information to ensure they are buying sufficient capacity and the appropriate hardware to meet their workload demand. There are many configuration options and choosing the best for each situation must be informed by accurate workload analysis and capacity planning.

This paper will focus on those operational procedures which are the responsibility of the customer. They include, but are not limited to:

- Planning
- Coordination
- Communication
- Testing
- Benchmarking
- Project management

Upgrades can include database software (major, minor, maintenance) and database hardware (upgrade, expansion, migration). This white paper will focus only on hardware upgrades. It will cover upgrades, expansions (including coexistence) and floor sweeps. The assumption is that software will remain static and sites have procured current site installation guides from their CSR and have reviewed them. We suggest that you perform software and hardware upgrades separately if possible to avoid too many moving parts at one time. This will help you identify causes of problems that may occur and is considered best practice for both customers and Teradata Customer Services.

Teradata upgrades are different now than in the past. Server technology is advancing at such a fast pace that upgrades are actually requiring fewer nodes and smaller footprints. This has obvious benefits, but there are potentially negative implications to be aware of.

These new servers are smaller and faster, and they are pushing considerably more AMPs per node. With fewer nodes and less parsing engines, you may have fewer overall sessions available. Depending on your workload, you may need plans to mitigate this, such as purchasing extra memory or PE-only nodes.

Make sure you take the necessary precautions to prevent load and user problems. For example, make sure your load users are not asking for maximum sessions on load utilities. They will get a session for

every AMP. This could be very bad if you have several utilities running, 2 or 3 times the number of AMPS and half the number of available sessions that you had previously.

Note to the reader. *This guide is provided as a courtesy of the Teradata Partners Service Focus Team (SFT). It comes as is; with no implied support. The procedures represent the combined learning of the customer members who contributed to it. The procedures are meant to be high level guides and do not represent step-by-step procedures. They are accurate to the best combined knowledge of the team, but should always be followed with caution and good judgment as they are not guaranteed to be free from flaws.*

Visit the Service Focus Team website to provide feedback on this whitepaper, download additional whitepapers, or learn more about the SFT.

<http://www.teradata-partners.com/About/ServiceFocusTeam>

2. The Upgrade Process for Teradata Hardware

Upgrade Process Overview Diagram

This diagram gives an overview of the upgrade process. The body of this document will describe each topic in more detail:

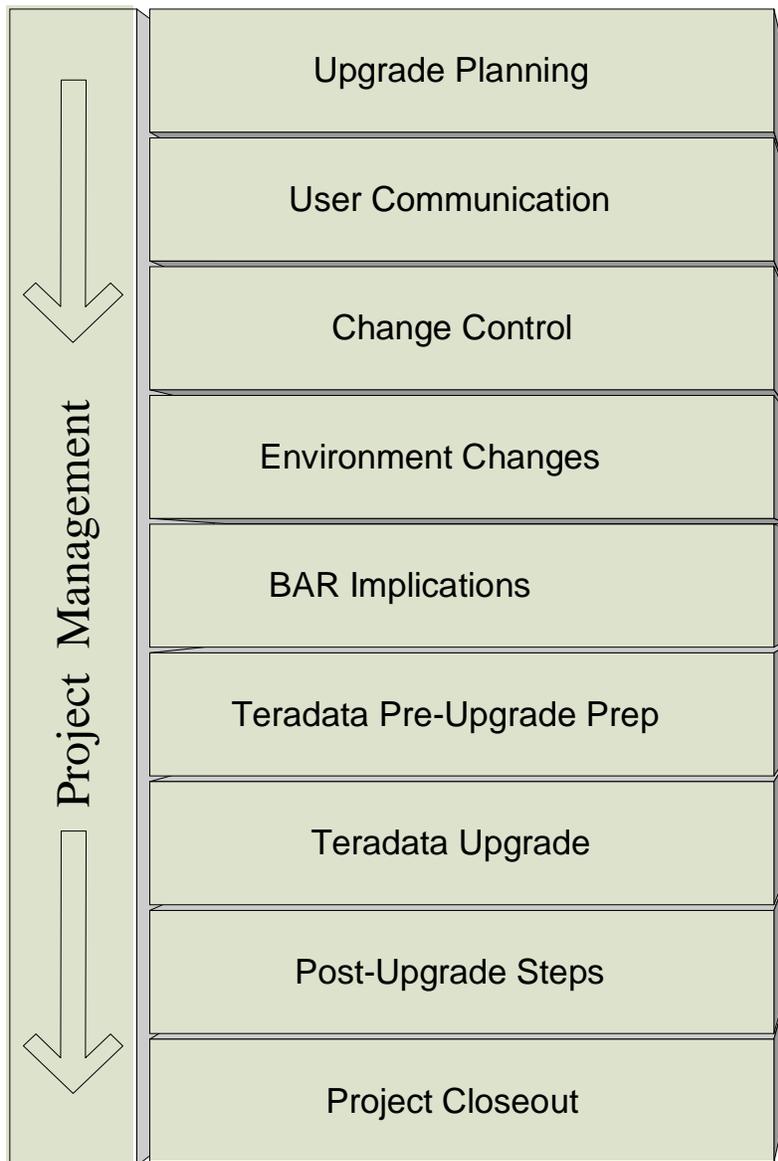


Figure 1: Upgrade Process Overview Diagram

3. Project Management

Any successful project requires strong and focused project management. A Teradata upgrade is a significant undertaking that is well worth treating as a formal project.

Assign a project manager to coordinate all the required tasks as a single project. This person will ensure that all involved parties communicate effectively, schedules are set and followed, and all tasks are completed successfully.

Create an upgrade team with DBAs, operations, developers and others to ensure cross-team participation and communication. Have a kick-off meeting to update everyone and then provide regular status updates to the members. You may want to include your Teradata Customer Service Representative (CSR) in the team meetings.

Some SFT member DBAs have found it helpful to maintain an upgrade plan template document that covers both hardware and software changes. It lists items to remember for specific circumstances. For example, when adding new nodes you would want to remember to request firewall rules changes, obtain IP addresses, make DNS entries, etc. This is a living document to which you would add new memory joggers learned during each new upgrade.

4. Upgrade Planning

An upgrade can be complex. Take the time up front to investigate the hardware changes before taking the plunge. The time invested up front will pay off later.

Plan for capacity

While Teradata can help, ultimately the customer has to decide how much hardware capacity his workload demands. The customer needs to collect this information and provide it to Teradata. The sales and site support team will use this information to provide hardware configuration options that can support the demand. There are many configuration options and choosing the best for each situation must be informed by accurate workload analysis and capacity planning. Also, see the BAR section of this document for information on how hardware configurations can affect backup and recovery.

Research environmental requirements

Data center environmental requirements can be quite complex. Your Teradata CSR can provide documentation for required site preparation that should be shared with your data center management team. Environmental requirements include, but are not limited to:

- Floor tiles, including both solid and perforated floor tile requirements. Work with Teradata to get a recommended layout of the new configuration, which will show where perforated floor tiles need to be located. This layout has to be approved by Teradata Engineering.
- New perforated tiles and electrical work may need to be ordered and scheduled. Allow time in your project plan for these.
- Verify with the CSR the type of electrical plugs needed for any new cabinets. This will need to be coordinated with the data center manager and/or electricians to make sure the proper electrical drops are available when you are ready to install. You may also need to consider the additional electrical load on your data center PDU's (Power Distribution Units).

- Find out if your data center requires the implementation of an EPO (Emergency Power Off) switch. These are sometimes required by fire departments that need to kill power to the entire data center from one switch in the event of a fire. This would need to be coordinated with Teradata, who has to wire the EPO in each node cabinet, and the data center manager/electrician, who extends the wires from the data center EPO to each cabinet.
- Your CSR may need to take temperature readings at various points of selected cabinets to check air flow and temperature readings both before and after the upgrade. Check to see if there is a temperature and humidity gauge in the data center and procedures for high temperature conditions.

Review network requirements

The customer's network administration team will need to coordinate with Teradata to provide the required network connectivity for the new nodes. Some of the things to address are:

- Network DNS (Domain Name Server) entries and IP addresses
- Verify security hardening has been applied to the new nodes
- Network hardware is available and configured including as cables, routers/switches, firewalls and rules, VPN (B2B or SSL certs), MVS networking (ESCON/FICON, Channel Extenders, type of ESCON connectors), etc.
- Investigate network load balancing. With the faster nodes, smaller footprints, dual active and hot standby nodes it has become more important to look at this aspect. Without some type of network load balancing a customer could end up with timeouts or a lack of network connectivity. Your Teradata CSR can help plan for this.

Pre-upgrade Benchmarking

If your management is concerned about the Return on Investment (ROI) of the upgrade, you'll want to identify queries you can use for before/after benchmarking. Types of SQL you may want to include are expedited queries, canary queries, heavy hitting queries, frequently run queries, ETL SQL and possibly some more complex queries that may have caused problems in the past. We recommend you identify at least 10-20 different queries.

One possible way to ensure consistency and security of your benchmark queries is to store them on the database as macros. Steps can be added to collect elapsed time, CPU and IO's and write them to a table.

Collect explain output, elapsed times and details for each query. DBQL is a great resource to capture this information. Be sure you have a backup of DBQL or whatever table you use to store this information so you are sure to have it after the upgrade.

If possible, run the queries when the system is quiet before and after the upgrade to accurately compare results. Note that run-time should be long enough to retrieve comparable runtime figures. Keep in mind they might run faster than usual on a quiet system.

You may want to include Fastload, Fastexport, Multiload, T-Pump, and archive and restore jobs in your benchmark suite.

Another important area you may need to consider is performance of a mixed workload. You will want to ensure that the various workloads will continue to perform as expected when the system is saturated. Creating a scenario to test this will help to prevent surprises after an upgrade.

A lot of factors can affect the performance of your benchmarks including networking issues, load on the system, missing indexes or join indexes, missing statistics, etc. Make sure all these are addressed before

beginning your benchmarks. If you run into slow benchmarks, check these things first before continuing your analysis or commenting to Teradata.

5. User Communication

You will want to keep your users well-informed regarding the upgrade timing and content. Here are some places you can gather information on hardware upgrades to pass on to your users:

- Partners Presentations. Pass them on as-is or pick-and-choose slides for repackaging.
- Implementation manuals. These are available on CD and on the Teradata website at <http://www.info.teradata.com>.
- If appropriate, the project plan as documented by the Teradata Operational Services project manager.

Ideas for user communications include:

- Publish upgrade information and schedule on your intranet website.
- Request users to drop unneeded objects from their personal databases. This is important for hardware expansions as this can reduce the run time of the Reconfig utility.
- Verify with users that all their critical data will be backed up if you don't already regularly back up user personal databases.

6. Change Control

The argument for formal change control is that system upgrades are managed in a controlled and successful manner such that there is minimal negative impact on the users of the system.

Main points to consider and potentially include in your change control procedures are:

- Be very clear about when changes are to take place and how long they are expected to take.
- Are there pre-requisites that affect the order of changes?
- Will your own staff carry out the changes? Will Teradata employees make any of the changes? Will other 3rd party employees make the changes?
- All changes must have a back-out mechanism. Remember to test that mechanism where possible.
- Agree well ahead of time who will work on the project and obtain commitment that they will be available.
- Make sure your change control includes a list of contact names and details for use during the time of change.
- Be sure to know who will test changes and be sure that they have a proper test plan.
- Consider what progress needs to be communicated to whom and at what points. Communication can make or break the project.
 - A Steering Committee may need a monthly report, whereas others will need the very latest information during the actual upgrade.
 - During the upgrade give timely (maybe every 4 hours) updates on the upgrade progress.

The key to successful change is planning and communication. Change control helps ensure all eventualities have been considered and communicated.

7. Environment Changes

Implementing hardware changes will likely result in physical changes and impact to various aspects of your environment. These include the data center, MVS systems, node OS environment, and the customer network. At a minimum, electrical power and air conditioning will be affected, and in many cases floor space and layout will be affected. It is important to understand how the planned hardware changes will affect your environment so that proper planning can occur. Some of the high points to consider are:

Data Center:

Be sure to involve your data center management early to review the planned changes, as many of these subtasks take time and funding to arrange. Some of the areas to address are:

- Air flow in the data center and how it is affected by the introduction of new equipment.
- Electricity – including power drops and different plug types for different cabinet types: node, disk array, channel only nodes, non-TPA nodes, appliance nodes (Viewpoint, etc.), EPO connections and BAR hardware
- Placement of floor tiles: perforated, non-perforated
- Temperature increase mitigation
- AWS: footprint, modem, network connections

MVS Environment:

- Chipids
- LPARs
- TDP
- TDP startup commands
- TDP memory cell settings
- Channel Extenders (CNT/Brocade)
- Request any new objects needed on mainframe
- Consider any increased capacity requirements over network channel

Node OS Environment:

- Add new IP addresses to /etc/hosts file
- Modify default parameters for new user creations
- Create new user id's on new nodes: admin, DBAs, Teradata support
- Add new DNS entries to internal files
- Harden Unix nodes as recommended by Teradata
- Create CSR admin directories and copy Unix utility/admin scripts
- Install SUDO if applicable for restricted root access

Network Environment:

- Request new IP addresses for new nodes from network group

- Request addition of new DNS names for new nodes
- Request addition of new nodes into Content Services Switch or other load balancer (if used)
- Provide CSR with IP addresses and subnet mask for any nodes being added.
- Provide CSR with Channel Unit Address (CUA) information for any channel-connected nodes being added.

8. BAR Implications

Do not ignore your backup and recovery (BAR) solution when upgrading your Teradata hardware. You should work with your Teradata sales and site support to ensure that your BAR infrastructure has sufficient capacity and enough bandwidth to backup your data within the available backup window.

There are three very common BAR configurations:

- LAN based BAR media servers with attached backup devices
- Mainframe using FICON/ESCON channels
- Teradata nodes using directly attached tape drives

There may be other approaches in the field, but these are the most common. In all cases as the number of Teradata nodes, amps or clusters change it is critical that you review and reconfigure your BAR configuration so that it can efficiently handle the new amp/cluster configuration.

At a bare minimum, when the cluster configuration changes you must make corresponding changes to all of your cluster backup scripts/jobs. If you are using NetVault or NetBackup with the legacy Teradata Extension, you'll need to change the number of clusters the software expects to match the cluster count on your Teradata server.

You may also want to increase your BAR parallelism to echo the increase in your Teradata server capacity. On channel attached systems, this could mean adding channel attachments (ESCON, FICON, etc) to more nodes. For BAR server configurations, it could mean adding additional BAR servers with backup devices.

For media server configuration all Teradata nodes and BAR servers should be configured on the BAR private network. Also, you should consider changing the BAR server host file on each media server to rebalance the "COP" mappings so they are evenly distributed across your Teradata nodes. Spreading the I/O evenly across the nodes helps ensure against network bottlenecks.

Speaking of I/O bottlenecks, as mentioned earlier in this document be aware that the data density per node has increased dramatically with the latest Teradata hardware generations like the 5550H. With recent floor sweeps it is not uncommon for the total data capacity to increase while the number of nodes decreases. That's great for cost reduction and "green" concerns, but can very negatively affect your BAR throughput. Basically you end up with fewer pipes to push all that data through.

Be sure to discuss this with your sales and site support team before signing the final contract to ensure that you are buying enough I/O bandwidth to keep your backups within your available backup window.

9. Pre-Upgrade Preparations with Teradata

The list of things to do in preparation is different depending on the type of upgrade being performed. Simple upgrades like upgrading CPUs or adding memory to existing nodes require less attention than do

expansions (where nodes are replaced or new nodes are added to an existing system) and floor sweeps (where an entire system is replaced with a new one).

Things to do for all upgrades include:

- Open change controls – Teradata & local.
- Ensure you have backups of DBS control settings, SysSecDefaults, ResUsage, TDWM, XCTL, PSF settings, etc. in case you need to do a full recovery if something goes wrong.
- If Teradata Customer Service system health tools have identified any preexisting issues that may impact the hardware upgrade, allow adequate time before the scheduled hardware upgrade to correct them. At the least, allow adequate time at the beginning of the hardware upgrade to correct the issues. While the goal is to minimize the downtime, it is important to go into one of these windows with stable hardware and OS.
- Before starting the upgrade window, inform the CSR if system hardening has been performed. Sometimes issues are encountered that are later found to be caused because of system hardening. If this information is known up front, it might help in the problem identification and resolution.
- If system hardening was performed, it may be necessary to enable services/ports that are needed for the tools (i.e. Teradata PUT) to operate properly.
- The PUT utility will log its actions while installing and configuring software. You may want to request these logs from your CSR and review them prior to deploying the new hardware to your users.

System expansions require new nodes to be added to an existing system, or swap out of old nodes for new ones. Usually new disk arrays are added as well, or disk drives are replaced with larger ones. Almost always the amp/cluster configuration of the system will change requiring Teradata to run the Reconfig utility to redistribute the existing data across the nodes. Things to address in this situation are:

- Discuss your burn-in requirements with your CSR. Teradata has a utility called Rescribe that will constantly exercise the new hardware for burn-in. Teradata's standard practice is to run Rescribe on new nodes and disk arrays for 72 hours, but some shops believe that is not enough time to ensure against hardware failures after deployment. Our recommendation is to run Rescribe for as long as you can prior to loading data onto the new nodes. Two to three weeks is not uncommon.
- Clean up any unnecessary data by deleting rows or dropping the tables. Look for things like work tables containing data of a temporary or transient nature. The more data on the system, the longer Reconfig will run.
- Review all tables using Fallback and remove that option if it isn't really required. Fallback will increase Reconfig time.
- Run scan disk and check table at level two within 30 days of the upgrade and correct any identified issues to be sure there are no pre-existing problems. Always use the Concurrent Mode option when running check table on an active system to reduce the chance for deadlocks.
- Ensure there is adequate free space available for Reconfig workspace. Depending on the amount of fallback tables in existence and the change in DBS clustering, Reconfig will require a significant amount of free space to distribute (i.e. temporary re-write to new location) data rows.
- Running packdisk is recommended to avoid mini cyl-packs during the Reconfig redistribution phase. To save time, the packdisk may be monitored from a second Ferret session running Showspace in a separate CNS window, and Packdisk aborted when adequate free cylinders is reached.
- Let the CSR know if there are PPI tables. The CSR can run a query to identify any PPI tables that might have an impact on the performance of the Reconfig. See Teradata knowledge article SD1000C4342 for more information.
- Run check table with PendingOps option immediately before the upgrade. If any tables are identified, take corrective action by dropping the tables or completing the utility operation.

- Drop any value-ordered indexes prior to the Reconfig, as they are not redistributed. Add a task to the project plan to recreate them after the Reconfig is completed.
- Run full backups of important data. All-amp backups are preferred over clustered backups if the amp configuration will be different after the upgrade. Remember that cluster backup datasets must be restored serially to a different amp configuration. That may require an unreasonably long restore window. Be sure to allow enough time in your schedule to complete the backups before the upgrade outage window.
- Clean any leftover spool space.
- If you use journaling, you may need to drop your journal tables prior to the upgrade.
- Disable DBQL and offload to history repository.
- Let your CSR know the location of User Defined Functions (UDFs) libraries before the outage window. This will help avoid delays because of trying to identify the location during the system expansion outage window. There are steps in the Teradata process to copy these libraries to the new nodes at the optimal point in the operation. Not performing this copy may result in restart loops when trying to bring up the database after the system expansion.
- Ensure that NTP (Network Time Protocol) is configured and running on the new nodes. It is very important to keep the time all the nodes synchronized. Linux is especially susceptible to time drift.
- Check the network interface cards (NICs) on the new nodes. Verify that the auto-negotiate, speed and duplex settings are correct for your network and consistent across all the nodes.

In a floor sweep an entire system is replaced with a new one. The usual approach is to have both systems running side-by-side so that data can be transferred using the NPARC utility. Many of the tasks listed above for expansions also apply to floor sweeps. For example you'll want to get rid of unnecessary data, set up NTP and run all the data integrity checks. But rather than run Reconfig, the data will be ported via NPARC. Some steps specific to floor sweeps are:

- Make sure your Data Center infrastructure is in place to house the new system. The upgrade steps really depend on whether you have the space to house two concurrent systems or not. This will determine how you handle the entire upgrade. Either way, you need to make sure you have floor space, power, air, etc. to support the new Teradata hardware, storage, BAR servers, AWS, etc.
- Provide the CSR with the desired time zone to be configured on the system. There are cases where the system is physically located in a different time zone than the users.
- If the system default names are not desired, provide the CSR with the desired node names so it can be set up during the installation, rather than later after the system has been put into production.

10. Teradata Upgrade

Now that the system is fully prepared, the actual upgrade may begin. While the hardware upgrade itself will typically be done by your Teradata CSR, these are last-minute functions the customer DBA may need to perform:

- Quiesce the system.
- Abort remaining user sessions.
- Backup DBC.
- Bring down any TDPs.
- Shutdown Teradata Manager server if any.
- Shutdown open systems backup application (NetBackup, NetVault).
- Disable unnecessary user ID's.
- Stop any job schedulers such as AutoSys, cron, etc.

- Run checktable Pending OP and resolve any identified issues.
- Ensure that all Hash, Join Indexes, Triggers and UDFs have been dropped.
- Turn the system over to the Teradata CSR to perform the hardware upgrade.

It may be required that the DBA performs a restore of the entire system from backup. If this is the case, allow time for an exceptionally long restore time. If you are doing a floor sweep or have a swing system, you can perform an NPARC to the new system. In that situation you may be able to do much of data transfer prior to the change date. Then you will only need to copy over the recent changes.

Be prepared for the Reconfig to take a long time. Some customers have experienced a Reconfig taking substantially longer than expected. Your CSR has a utility that estimates how long the Reconfig should take, but some object types (especially PPI tables) can cause the Reconfig take longer than predicted.

At the appropriate point during the process, the CSR will compile any user defined functions (UDFs) and stored procedures you may have to deploy them to the new nodes. If there is a problem doing that (missing source code, etc) the CSR is supposed to let you know. We recommend that you follow up with your CSR to ensure the compiles completed successfully.

11. Post-Upgrade Steps

After the CSR has completed and validated the upgrade, he or she will turn the system back over to the DBA/System Administrator. The DBA/SA would then perform a sequence of functions, including:

- Check AMP configuration, and related aspects like spool, DBS control settings, PE's, etc. If running coexistence, be aware of AMP configuration and imbalances, and make sure channel cards are in the newest nodes or channel only nodes (avoid TPA nodes).
- Any new monitoring (TVI, Enterprise Monitoring Tools, etc.)
- Reestablish TASM/PSF CPU-based limits.
- If you have hot spare nodes, test fail over and make sure this works with 'round robin' node selection.
- Run checktable DBC only at level 3. This should only take a few minutes, depending on the number of objects.
- Run packdisk.
- Run scandisk and checktable to verify the rest of the file system is clean. See the Teradata Utilities Volume 1 A-F manual for help choosing the appropriate check level.
- Check/reset TDWM, DBScontrol, XCTL and PSF settings.
- Restart DBQL, Access logging, System Security, Priority scheduler, etc. as before.
- Restart the Teradata Manager server.
- Recreate Hash, Join Indexes, Triggers and UDFs that were dropped prior to the upgrade.
- Recreate journal tables (if any).
- Make sure the default date format is set correctly. This often gets changed.
- Recompile and test any user stored procedures.
- Validate statistics are collected as should be on DBC and user tables. Recollect if necessary.
- May need to adjust user spool settings if number of amps changed.
- DBA regression testing. Run test program to validate all settings, connectivity, programs, DBC tables, procedures, etc.
- Test connectivity of applications such as SQL Assistant, BTEQ, etc. to the PEs on the new nodes. Start/confirm TDP's and enable logons.

- Restart job schedulers such as AutoSys, cron, etc.
- Restart open systems backup application.
- Notify customers of system availability.
- Sign off on user acceptance of upgrade with Teradata.
- Close Change Controls – Teradata & Local.

At this point you could run your post-upgrade benchmarks and compare them to your pre-upgrade benchmarks. Publish results.

12. Project Closeout

Shortly after the upgrade is complete and users are happily using the system you will want to have a project wrap-up meeting. The project manager should lead a “lessons learned” session and document what went well and what could be done better next time. This information will be invaluable for planning the next major upgrade.

If you are maintaining an upgrade plan template document (see the Project Management section above), this would be a very good time to update it with anything new you have learned.

13. Closing

The SFT is a committee of the PARTNERS User Group that works closely with Teradata on issues related to support services and other areas that fall beyond the scope of product enhancements. Members represent the concerns of Teradata customers by serving as catalysts for service improvements, providing ongoing feedback to Teradata and the PARTNERS Steering Committee.

The members of the SFT sincerely hope that you found this document useful. Please let us know of any comments or suggestions you may have. See you at Partners!

You can find our contact information and more about the SFT at our website:

<http://www.teradata-partners.com/partners-user-group/service-focus-team>