Task 4

1st query is reasonably fast because it has only one condition to satisfy, the query was fast because it wasn't going through a lot of records and it has no second condition to meet.

2nd query was reasonably slow because it has two conditions to satisfy. If the 1st condition is met, the query has to check again if the second condition is met, thus slowing down the query.

Task 3

Modern disk systems are generally reliable. However, every now and then things do go wrong, and data is not written successfully to disk even though the disk subsystem and the operating system believes that it has. These types of events can cause widespread data corruption in a database environment, and they can be difficult to detect. Indeed, it may not be until the next access of that data that any inconsistency is noticed, allowing the problem to spread through the entire database and possibly through backups and disaster recovery sites as well.

Oracle Database has comprehensive built-in checks to detect and repair data corruptions. Using a single parameter set to a desired protection level, Oracle Database 11g , can detect corruptions in data and redo blocks using checksum validation, detect data block corruptions using semantic checks, and detect writes acknowledged but actually lost by the I/O subsystem. Specific technologies also provide additional validation: Oracle Recovery Manager can be used to validate data blocks while doing backup and recovery, and Oracle Automatic Storage Management can be used to recover corrupted blocks on a disk by using the valid blocks available on the mirrored disks. Plus, Oracle Data Guard can be used to ensure that the standby database is isolated from all data corruptions at the primary database.

Oracle Exadata Database Machine also prevents corruptions from being written to disk by incorporating the hardware assisted resilient data (HARD) technology in its software. HARD uses block checking, in which the storage subsystem validates the Oracle block contents, preventing corrupted data from being written to disk. HARD checks with Oracle Exadata operate completely transparently and no parameters need to be set for this purpose at the database or storage tier.

Protect Against Data Loss with Oracle Automatic Storage Management Mirroring

Many IT professionals use disk-mirroring techniques such as RAID 5 or RAID 10 to keep multiple copies of data on different disks in their storage arrays, protecting against data loss from an individual disk failure. Oracle Database 11g works very well with all major disk-mirroring capabilities. However, these disk-mirroring capabilities can often add additional expense to the storage environment being used.

Oracle Automatic Storage Management can be used to not only stripe data across multiple disks in a storage array, but to also mirror that data across disks as well.

Redundancy for disk groups can be either normal (the default—files are two-way mirrored, requiring at least two failure groups), or high, which provides a higher degree of protection using three-way mirroring (requiring at least three failure groups).

Oracle Automatic Storage Management uses a unique mirroring algorithm that mirrors extents. When Oracle Automatic Storage Management allocates a primary extent of a file to one disk in a failure group, it allocates a mirror copy of that extent to another disk in another failure group, ensuring that a primary extent and its mirror copy never reside in the same failure group. Unlike other volume managers, Oracle Automatic Storage Management has no concept of a primary disk or a mirrored disk; a disk group only requires spare capacity, so a hot spare disk is unnecessary.

When a block is read from disk, it is always read from the primary extent, unless the primary extent cannot be read, in which case it will be read from the secondary extent. When a block is to be written to a file, each extent in the extent set is written in parallel.

In the event of a disk failure in a failure group (which will induce a rebalance), the contents (data extents) of the failed disk are reconstructed using the redundant copies of the extents from partner disks. If the database instance needs to access an extent whose primary extent was on the failed disk, then the database will read the mirror copy from the appropriate disk. After the rebalance is complete and the disk contents are fully reconstructed, the database instance returns to reading primary copies only.

Efficiently Backup and Restore Data

Although storage mirroring provides an important element of data protection, all databases should be regularly backed up. However, as databases grow in size, new optimized techniques are required to constrain both the time required to backup the database and the time required to restore and recover (if necessary).

Traditionally, data has been backed up to streaming devices such as tapes. Oracle Database 11g supports many tape backup and vaulting environments through Oracle‘s backup solutions program, a cooperative program designed to facilitate tighter integration between Oracle's backup products and those of third-party media management vendors.

Oracle Secure Backup

In addition to supporting third-party backup products, Oracle provides a tape backup solution. Oracle Secure Backup provides centralized tape backup management, ensuring high-performance data protection of both file system and Oracle Database data in UNIX, Linux, Windows, and network-attached storage environments.

Oracle Secure Backup is a complete tape backup solution that is fully integrated with the Oracle Database backup utility, Oracle Recovery Manager. This tight integration means that Oracle Secure Backup only backs up currently used blocks, eliminating backup of committed undo, which helps to increase backup performance by 25 to 40 percent over comparable products. In addition, Oracle Secure Backup offers backup encryption and key management, ensuring that any confidential information stored on backup tapes sent offsite are protected.

Fast Recovery Areas and Incremental Backups

One of the challenges of a tape-based backup strategy is that writing backups to tape and subsequent restores can take a long time. More and more organizations are utilizing low-cost disks as their preferred storage for Oracle Database backups. One of the advantages of this approach is that random I/O can be performed on the backup images stored on disks. Oracle Database 11g takes full advantage of this capability with disk-based backup and recovery technologies.

Database administrators can define a disk-based fast recovery area for Oracle Databases. This is a group of disks typically separate from the storage array used for the database environment. The fast recovery area is dedicated for database backup images and fully managed by Oracle Recovery Manager, and it can take advantage of Oracle Automatic Storage Management for backup striping and the like. Once a fast recovery area is set up, Oracle Databases will automatically back themselves up during predefined backup windows. More frequent backup periods can be defined, along with the timing and duration of the backup window. If additional space in the fast recovery area is needed, Oracle Recovery Manager will automatically delete files that are obsolete or have already been backed up to tape.

For large databases, backing up the entire database every 24 hours may take a long time. Instead, many IT organizations enable incremental change tracking on their Oracle Databases. During the 24-hour period since the last backup, a record is kept of the data blocks in Oracle Database that have changed during OLTP operations. During the backup window, only the changed blocks are stored in the fast recovery area and can be used as an incremental backup against a complete backup image. This means that the backup window itself is only ever a function of the number of changes made within a 24-hour period, independent of the size of the actual database.

However, too many incremental backups can extend the time it takes to restore and recover the database if required. To this end, Oracle Recovery Manager incrementally updates the complete database backup image with the incremental changes from the previous 24-hour period, eliminating the need to apply multiple incremental backups on recovery. This reduces overall recovery time in addition to reducing the need to make full backups.

Backup images in the fast recovery area can be stored to tape using Oracle Secure Backup or other tape backup solutions. Such backups can also be compressed to further save time and space.

Read-Only Tables and Tablespaces

Read-only tables and tablespaces can reduce backup and restore times. Large tables that have been partitioned using a lifecycle management strategy can have older data partitions put into read-only format, so transactions cannot change the data stored within these partitions. These partitions can then also be compressed and placed on low-cost storage tiers. Because they are read-only, Oracle Recovery Manager knows that they do not need to be backed up beyond the initial backup; further reducing the requirements for both backup and restore operations.

Using low-cost disk storage for the fast recovery area and an incremental backup strategy provides an alternative to expensive snapshot technologies deployed at the storage array level. Oracle Exadata Database Machine makes use of these data services and comes preconfigured with a fast recovery area for the databases that are built on them.

Protect Against Data Loss Caused by Human Error

Most data loss is caused by human error, and not the failure of a disk or a storage array. Database administrators (DBAs) make mistakes—they log onto development databases to clean up tables and indexes only to find that they logged onto the production environment by mistake and dropped critical tables and indexes that were in use at the time.

When these problems occur, the traditional approach has been to revert to a backup. The production system is stopped and a point-in-time recovery is performed to restore the data back to the point just before the error occurred. There are many problems with this approach—additional storage is required, restore and recovery can take a long time while the production system is unavailable, and any good transactions made from the time the error occurred until it was later discovered are lost.

To overcome this, Oracle Database 11g provides unique flashback capabilities. Flashback allows operations that were inadvertently performed online to be undone online. For example, if a DBA drops a table or an index, the table and index are marked as unavailable in the data dictionary and are no longer used by the application. However, the actual data extents that were present in the table or index are kept on disk. If at a later stage the DBA identifies that the table or index was dropped by mistake, he can simply undo the drop operation and the table or index is immediately made available again to the application.

Similar steps can be taken if a transaction invalidly changes one or more rows in a table. Flashback query operations allow the DBA to see the earlier versions of the rows and identify the transactions that caused the error. Then the offending transactions can be flashbacked online, so all changes caused by the transactions are immediately undone. Alternatively, all changes made to one or more tables since a specified period of time can also be undone online.

If the entire database has become logically corrupt due to a large number changes, the entire database can be flashed back in time, rather like running a video in reverse. While flashing back, the database is not in online operation, so it is far quicker to unwind the database this way than restoring a database and doing a point-in-time recovery. In addition, the flashback operation can be performed multiple times, allowing incremental flashback and roll forward to the exact moment in time the error occurred.

Flashback transactions and flashback table operations use the existing undo information that is collected by Oracle Database during normal operations. Most Oracle DBAs will size the collection of this undo information to give themselves a 24-hour window in which to catch and undo any human error. The ability to flashback the entire database requires additional information to be collected. This additional information is managed as part of the fast recovery area and is also typically sized to provide a 24-hour window.

This flashback capability supplants the need for additional snapshots at the storage level, further reducing storage costs and providing a much finer granularity of operation. It also provides the ability to undo any individual transaction to any point in time.

Protect Against Data Loss Caused by Disaster

The techniques just outlined protect against data loss within the datacenter. However, the datacenter itself can be lost through fire, earthquake, other disaster, or even something as mundane as power loss. To fully protect themselves against data loss, many IT organizations are investing in secondary datacenters with standby databases, which can be synchronized with the changes being made in the production environment.

The traditional method of synchronization requires expensive, remotely mirrored storage solutions. The storage replicates every write performed on the production system to the standby system. This means that expensive, high-bandwidth networks are required between datacenters, incurring additional cost and limiting the distances that can exist between the datacenters.

The remotely mirrored standby solution is only used in the event of the loss of the production datacenter. An idle resource decreases the value of the investment made in the standby environment, and administrators are often reluctant to failover to the standby environment because they are not familiar with its day-to-day operation.

Remote mirrored storage also introduces another layer in the infrastructure where problems can occur, especially concerning the propagation of data corruption between the production and standby sites. Oracle Database works well with remotely mirrored storage solutions, but also provides a built-in standby solution, Oracle Data Guard, which addresses the problems of remote mirrored storage.

Oracle Data Guard only transmits the writes from the production databases to the standby databases. This can mean a 1/7th reduction on the network of the volume, and up to 1/27th of

I/O operations, allowing a much smaller-capacity network to be utilized between the datacenters,

- Have your DBAs verify that the backup is succeeding regularly, preferably using a script that notifies them if there's an issue.
- Maintain a backup to your backup. DBAs should always use at least two backup methods. A common technique is to use those old-fashioned exports as a backup to the online backups.
- Resource test recoveries as often as is practical. An early sign that your DBA team is either overworked or not prioritizing correctly is having a quarter go by without a test recovery. Test recoveries confirm that your backup strategy is on track, while allowing your team to practice recovery activities so they can handle them effectively when the time comes.

- Inculcate a "practice, practice, practice" mentality throughout the organization. DBAs need to rehearse activities in the safe sandbox of a test environment that's designed to closely mimic the behavior of the production system. The organization needs to allow the time and money for them to do so.
- Pair inexperienced DBAs with senior ones whenever possible—or take them under your own wing. New DBAs tend to be fearless, but learning from someone else's experience can help instill some much needed paranoia.

- Review the plans for everything. It's amazing how often DBAs say, "I've done that a hundred times, I don't need a plan." If they're heading into execution mode, they absolutely need a plan.

- Install availability and performance monitoring systems so that issues are identified and resolved before they cause service-affecting failures.
- Avoid post-release software issues by working with developers and testers to ensure that all production-ready software is stable and high-performance.

- Require that your DBAs maintain a comprehensive documentation library and activity diary, including a significant level of rationale, syntax, and workflow detail.
- Provide your team with groupware on your intranet so that these documents become searchable in an emergency.
- Enforce the discipline of documentation and check it periodically. Ask your DBAs: When was this tablespace created, by whom, and with what SQL? What tasks were performed on a particular day? If they can't answer quickly, you'll know they've gone

- Never upgrade your hardware infrastructure without first exhausting all tuning opportunities. Remember, ten years ago enormous enterprises were run on servers one-tenth the capacity—all thanks to necessity and skill.
- Never consent to using advanced or new features until you're well aware of the ongoing maintenance commitment and resulting costs.
- Watch out for DBA support software that presents friendly GUI interfaces for difficult tasks. This type of interface allows a beginner DBA to act as an intermediate DBA under certain circumstances, but simultaneously prevents that beginner from learning the actual skills behind the tasks. Moreover, these tools tend to hide real risks from the DBA, making potentially damaging activities as easy as point-and-click.