

Day - 1

• A A A

• A A

• A

A A A A A A A A A A

A AA A A

A

A A A

- Usage of initdb command
- Concept and structure of database cluster
- Template database

• A A A A A

A A A A A A A A A A A A A A

A A A A A A A

A

A A A

- Usage of database management commands

• A A

A A A A A A A A A A

A A

A

A A A

- About postgresql.conf
- How to set pg_hba.conf
- Usage of SET / SHOWA

• C A A A

A A A A A A A A A A

A

A

A A A

- Usage of various backup commands
- File system level backup and restore
- Concepts and procedures of Point-in-Time Recovery (PITR)
- Transaction log (WAL) and WAL archive
- pg_start_backup () / pg_stop_backup ()
- Usage of COPY statement (SQL) and \copy command (psql)

• C A A A

A A A A A A A A A A

A A

A

A A A

- Starting and stopping PostgreSQL
- Concept of database role / user
- Adding, deleting and modifying database role / user
- Purpose and usage of VACUUM and ANALYZE
- Autovacuum concepts and behavior
- System information functions
- Information schema and system catalogs
- Table privileges (GRANT / REVOKE)

Day - 2

• **A A**
A A A A A A A A A A A A A

A

A A A

- SELECT statement
- INSERT statement
- UPDATE statement
- DELETE statement
- Datatypes
- Table definition
- Index
- View
- Materialized view
- Rule
- Trigger
- Sequence
- Schema
- Table space
- Partition
- Function definition / Procedure definition
- PL / pgSQL

• **A A A**
A A A A A A A A A A A A A

A

A A A

- Aggregate functions
- Mathematical functions
- Operators
- String functions
- String operators
- Date/Time functions

• **A**
A A A A A A A A A A A A A

A

A A A

- Transaction syntax
- Transaction isolation level (read committed, repeatable read, serializable)
- LOCK statement
- Row-level locks and table locks
- Deadlock



A A

A

A

A

A

A

A

A A

A

A

A

A

A

A

A

A

A

A

A

A

A

A

A

A

○ Estimation of table / index size

○ Security

○ Encryption of communication route (SSL)

○ Data encryption

○ Client authentication

○ Audit log

○ Bootstrap Windows Nodes

○ Running the Cookbooks

• A A A A A A A A A

A

A A A

- Backup, PITR
- VACUUM, ANALYZE, REINDEX
- Autovacuum
- Checkpoint
- Server log management
- Disk space monitoring
- Difference between automatic VACUUM and manual VACUUM / ANALYZE

• A A A A A A A A A A

A

A A A

- Architecture of database cluster
- Database processes
- Physical structure of data and data files

• A A A A A A A A A A A A A A A

A

A A A A

A A A

- Overview of streaming replication feature and logical replication feature
- Synchronous and asynchronous replication
- Settings of postgresql.conf and recovery.conf
- Definition of publications and subscriptions

• A AA A A A A A A A A A A

A

A A A A A A A

A A A

- pg_locks
- pg_stat_activity, pg_stat_database
- Row-level statistics, ex. pg_stat_all_tables
- Block level statistics, ex. pg_statio_all_tables

• A A A A A A A A A A A A

A

A A A A A

A A A

- pg_class
- pg_stats
- Real file and path of table index
- Statistics and parameters used during execution planning

• A A
A A A A A A A A A A A A A

A

A A A

- EXPLAIN / EXPLAIN ANALYZE output
- Planning type
- Tuning from EXPLAIN
- Type of join (Nested Loop, Hash, Merge) and performance characteristics
- Execution plan corresponding to SQL syntax (JOIN / GROUP BY / ORDER BY / LIMIT)
- Execution plan corresponding to SQL with aggregate function (sum / count)
- Execution plan corresponding to SQL for partition
- Execution plan corresponding to parallel query
- Execution plan corresponding to SQL of window function (row_number / rank etc.)

• A A A A A A A A A A
A A A A A A A A A A

A

A A A

- Detecting slow queries
- Analysis using contrib tools
- Performance degradation factors (resource exhaustion, lock contention)

• A A A A
A A A A A A A A A A A A

A

A A A

A A A

- RESOURCE USAGE
- WRITE AHEAD LOG
- QUERY TUNING
- RUNTIME STATISTICS
- LOCK MANAGEMENT
- Light and heavy locks

• A A A A A A A A A A
A A A A A A A A A A

A

A A A

- Tuning parameters
- Tuning the execution plan
- Tuning SQL
- Tuning the table configuration
- Distributing disk I/O
- How to reflect parameters (action required for parameter activation)
- How indexes affect SQL performance
- Index Only Scan and Visibility Map

- **A A A**
A A A A A A A A A A A A
A A AA A A A A A A A A

A

- **A A A**
 - Transaction log recovery
 - System table index recovery
 - Developer options
 - Real file and path of table index
 - Relfilenode and OID
 - Index corruption and REINDEX recovery
 - Checksum corruption detection and recovery
 - Transaction ID rounding error

- **A A A A A A A A A A A A**
A A A A A A A A A A A A

A

- **A A A**
 - Streaming replication and logical replication
 - Error messages in log files
 - Processing that is/isn't propagated to the standby
 - How to restart primary PostgreSQL at stop/failure
 - How to restart standby PostgreSQL at stop/failure
 - Processing that is/isn't propagated to logical replication subscribers
 - Conflicts at Logical Replication Subscribers

- **AA A**
 - Orafce
 - Pg Pool II
 - Oracle_fdw
 - Pg_statsinfo
 - Pg_hint_plan
 - Pg_dbms_stats
 - Pg_repack
 - Pg_rman
 - pgBadger
 - pg_bigm
 - PostgreSQL JDBC driver
 - psqLODBC
 - Npgsql