Bonjour tout le monde

Installation
Tour de table
Présentation
Démarrage

Passw@rd
Binaires d'install dans /install
cible dans /app.

Oracle
1 Binaire
|
INSTANCE
|
1 jeu de fichiers de données } La BDD

MS SQL
PG
MySql · MariaDb
etc...

1 binaire = 1 version

/

1 Instance
= 1 binaire
en mémoire

/

1 jeu de fichiers de base
/master, tempdb, ...

4 BDD système
+
0_n BDD utilisateur

Chaque instance a ses ports TCP dédiés

---

80 → Microsoft → pas de SGBDR

Oracle v1 → 1979

Sybase → Sybase Sql Server

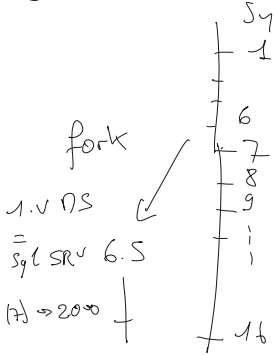

→ fin de Contrat

master → BDD Système !

tempdb → opération temporaires
→ objets temporaires
Create table #ma table
→ tri, jointures, ...
→ CRITIQUE POUR LA PERF.

---

useradd sybase
mkdir -p /var/ase/i1

SERVER NAME → ip du Serveur !
admin password → Sa / ...
(System administrator)

chown sybase /var/ase/i1
chgrp

mkdir /var/ase/i1/data
ajouter "Linux" dans /etc/hosts

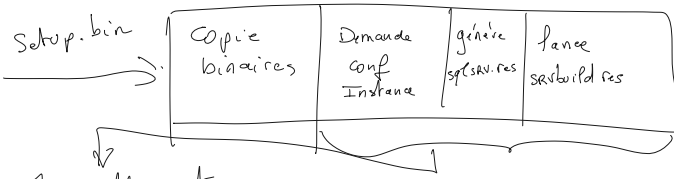

Manuellement

(1) → fichier .res → SRVbuildres

(2) → SRVbuildpl → erreur lib Xm.so.4

manque motif (yum install motif)

Pour lancer `pl/instance.pl`

Dans `<dossier d'instance>/ASE-16_0/install`

Lancer RUN_LINUX

host

`<INSTANCES>/ASE-16_0/LINUX.cfg`

Paramètre de lancement du Serveur : Quantité de mémoire

Lancer le serveur :

SOURCE <bindir>/SYBASE.sh

[export LANG=en]

<datadir>/ASE-16_0/install/RUN_linux.

Connection :

isql -S ASE:5000 -U sa

le nom du serveur à la création de l'instance, et doit aussi répondre au ping

---

[Nom de Serveur].[BDD].[user].Objet

master.. sysconfigures

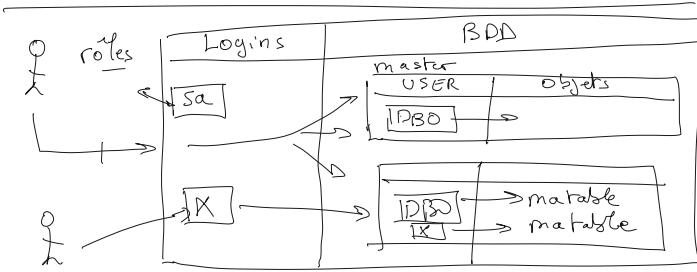

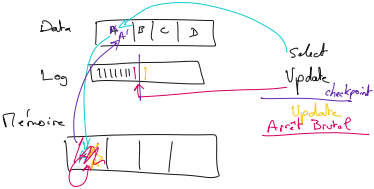

disk init name='mabase', physname ='/var/bdd/i1/data/mabase.dat', size='100M'
disk init name='mabaselog', physname ='/var/bdd/i1/data/mabase.log', size='100M'
create database mabase on mabase="10M" log on mabaselog="10M"
dbcc checktable(syslogs)

Création de login :
create login student with password 'password'

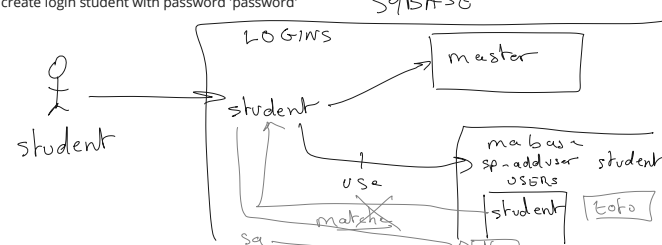

Création d'un user :
use master
create login student with password 'password'
create login titi with password 'password'
use mabase

sp_adduser student -> **le login et le user sont student)**
sp_adduser titi, toto -> **le login est titi, le user associé dans la base est toto**

sp_helpuser student
sp_helpgroup
create table test (id int)

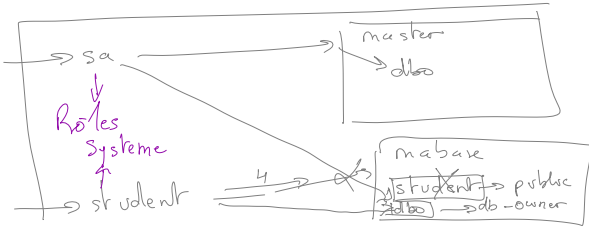

sp_dropuser student
sp_addalias student,dbo

par base :
sp_addgroup grpname (concept)

sp_changegroup management, jim (affecter un user à un autre groupe)

grant insert, delete on titles to mary (user) , sales (groupe)

exemple :
use mabase
sp_addgroup management
GO
use master
GO
create login jim with password 'password'
GO
sp_helpgroup
sp_helpuser
sp_adduser jim, jim, management
grant create table to management

Tuning :

Comprendre l'arbre de requêtes
Le contrôle du plan d'execution
Table avec et sans clef primaire
Index secondaire sur table avec, et sans PK
Paramètres de tuning, hints, ...

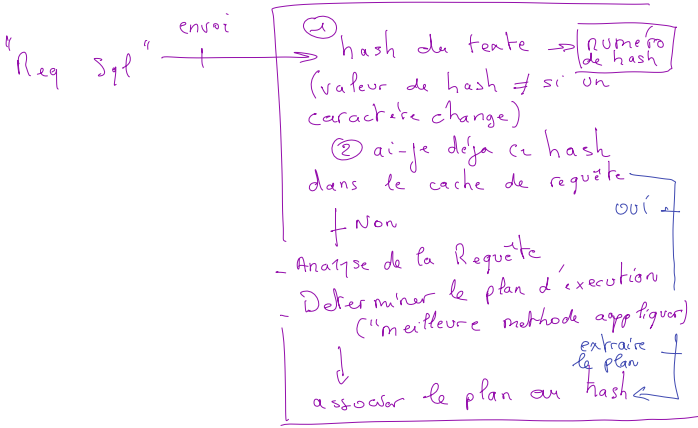

Influences sur le plan :
→ Répartition des données - objectif, plus ↓ de pages possibles.

Pas de PK

Insert → l'ordre physique des lignes
→ l'ordre d'insertion

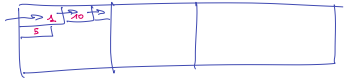

SYBASE autorise

0..1 Index cluster

0..n Index non cluster

cluster = tri au niveau données

non cluster = index (secondaire) avec ses pages d'index

PRIMARY KEY = Index cluster Unique

Taille d'index, pages

1, 10, 5 ⇒ 1 5 10

(déplacer les lignes)

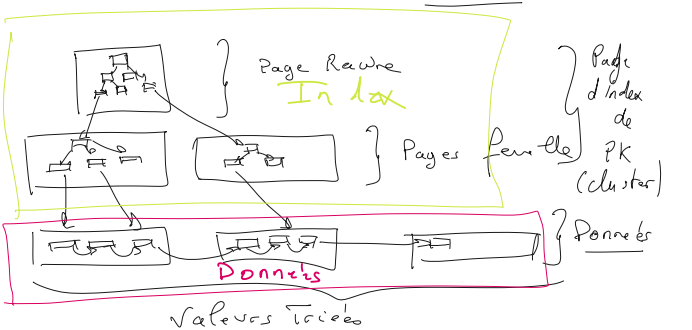

Index Secondaire / PK	avec	sans
avec	feuille de l'index = valeur de PK	Scan / Recherche via PK uniquement
Sans	feuille d'index Référence la Page de data	Scan uniquement

Table

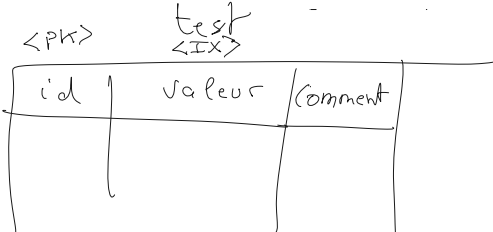

Select * from — where id= ⇒ PK

Select valeur from — where valeur= ⇒ IX

" id from — " valeur = ⇒ IX

" Comment from — " " = ⇒ IX Table

Index courrant = index qui inclut d'autres colonnes en + de la PK

isql -S ASE:5000 -U sa
create database test
GO

Executer une restore :
isql -S ASE:5000 -U sa -D test -i /app/ase16/jConnect-16_0/sample2/pubs2_any.sql

Lister les tables :
select * from sysobjects where type='U'

détail sur une table :
sp_help <nom de table>

select * from sales s inner join stores st on s.stor_id=st.stor_id

isql -S ASE:5000 -U sa -i /app/ase16/ASE-16_0/scripts/installpubs2

Dupliquer la BDD:

- Effectuez un Dump
  - Creez les devices pour la restore
  - creez la BDD de destination
  - Chargez : load -

Options importantes de dump database:

noinit (init):

Défaut — init → réinitialise le contenu

```
1> dump database pubs2 to '/var/bdd/pubs2.dump'
2> go
Backup Server: 4.171.1.1: The current value of 'reserved pages threshold' is
85%.
Backup Server: 4.171.1.2: The current value of 'allocated pages threshold' is
40%.
Backup Server: 4.171.1.5: The current value of 'parallel scan' is 2.
Backup Server session id is: 5. Use this value when executing the
'sp_volchanged' system stored procedure after fulfilling any volume change
request from the Backup Server.
Backup Server: 4.41.1.1: Creating new disk file /var/bdd/pubs2.dump.
Backup Server: 6.28.1.1: Dumpfile name 'pubs21920609E50  ' section number 1
mounted on disk file '/var/bdd/pubs2.dump'
Backup Server: 4.188.1.1: Database pubs2: 796 kilobytes (40%) DUMPED.
Backup Server: 4.188.1.1: Database pubs2: 1394 kilobytes (100%) DUMPED.
Backup Server: 3.43.1.1: Dump phase number 1 completed.
Backup Server: 3.43.1.1: Dump phase number 2 completed.
Backup Server: 3.43.1.1: Dump phase number 3 completed.
Backup Server: 4.188.1.1: Database pubs2: 1408 kilobytes (100%) DUMPED.
Backup Server: 3.42.1.1: DUMP is complete (database pubs2).
```

```
1> use master
2> go
1> disk init name="pubs20", physname="/var/bdd/i1/data/pubs20.dat", size="100M"
2> go
1> disk init name="pubs20log", physname="/var/bdd/i1/data/pubs20log.dat", size="100M"
2> go
```

```
create database pubs20 on pubs20="20M" log on pubs20log="20M"
```

```
1> load database pubs20 from '/var/bdd/pubs2.dump'
2> go
Backup Server session id is: 11. Use this value when executing the
'sp_volchanged' system stored procedure after fulfilling any volume change
request from the Backup Server.
Backup Server: 6.28.1.1: Dumpfile name 'pubs21920609E50  ' section number 1
mounted on disk file '/var/bdd/pubs2.dump'
Backup Server: 4.188.1.1: Database pubs20: 2820 kilobytes (6%) LOADED.
Backup Server: 4.188.1.1: Database pubs20: 7174 kilobytes (17%) LOADED.
Backup Server: 4.188.1.1: Database pubs20: 7188 kilobytes (100%) LOADED.
Backup Server: 3.42.1.1: LOAD is complete (database pubs20).
All dumped pages have been loaded. ASE is now clearing pages above page 1792,
which were not present in the database just loaded.
ASE has finished clearing database pages.
Started estimating recovery log boundaries for database 'pubs20'.
Database 'pubs20', checkpoint=(1334, 37), first=(1334, 37), last=(1335, 1).
Completed estimating recovery log boundaries for database 'pubs20'.
Started ANALYSIS pass for database 'pubs20'.
Completed ANALYSIS pass for database 'pubs20'.
Started REDO pass for database 'pubs20'. The total number of log records to
process is 5.
Redo pass of recovery has processed 1 committed and 0 aborted transactions.
Completed REDO pass for database 'pubs20'.
Use the ONLINE DATABASE command to bring this database online; ASE will not
bring it online automatically.
1> use pubs20
2> go
Msg 950, Level 14, State 1:
Server 'LINUX', Line 1:
Database 'pubs20' is currently offline. Please wait and try your command again
later.
1> online database pubs20
2> go
Started estimating recovery log boundaries for database 'pubs20'.
Database 'pubs20', checkpoint=(1334, 37), first=(1334, 37), last=(1335, 1).
Completed estimating recovery log boundaries for database 'pubs20'.
Started ANALYSIS pass for database 'pubs20'.
Completed ANALYSIS pass for database 'pubs20'.
Recovery of database 'pubs20' will undo incomplete nested top actions.
Database 'pubs20' is now online.
1> use pubs20
2> select count(*) from sales
3> go
```

Effectuer un backup des transaction d'une base :
Outil de recherche dans les pdf en ligne de SAP :
https://help.sap.com/viewer/product/SAP_ASE/16.0.3.7/en-US

passer une BDD compatible avec le dump transaction :

1) s'assurer d'avoir un device pour les logs (disk init)
2) ajouter ce device à la base
3) migrer vers l'utilisation de ce device :
 a) sp_dboption pubs20, 'single user', true
 b) on indique le device a utiliser pour les logs :
    sp_logdevice pubs20, pubs20log
 c) on la repasse en multi user :
1> sp_dboption pubs20, 'single user', false
2> go
Database option 'single user' turned OFF for database 'pubs20'.
Running CHECKPOINT on database 'pubs20' for option 'single user'
to take effect.
(return status = 0)
1> checkpoint pubs20
2> go
IL FAUT EXECUTER UN CERTAIN NOMBRE DE REQUETES POUR QUE TOUT LE LOG PASSE SUR L'AUTRE DEVICE

Créez une base de données avec du contenu (pubs2 par ex) avec data et log séparés
create table sales (ord_id varchar(50), ord_num varchar(50), date datetime)
insert into sales select * from pubs2..sales
dump database
update ou delete sur test..sales
dump tran

-> restorer la base dans une autre, en full, et en full+transaction

```
drop procedure maproc
go
create procedure maproc( @p1 int )
as
declare
   @mavar int
begin
   select @mavar=2
   select * from sysobjects
   select @@rowcount*@mavar
end

drop function mafunc
create function mafunc returns int as return 1

maproc 10

select dbo.mafunc() -- obligatoire de spécifier le propriétaire sous Sybase
```