



六节课快速上手Greenplum

第一课

Greenplum 介绍、安装与部署

Greenplum高级解决方案架构师:李兴欣



目录



GREENPLUM
DATABASE®

腾讯云大学

第一节 Greenplum介绍

第二节 Greenplum市场地位

第三节 Greenplum架构设计

第四节 Greenplum机器选型

第五节 Greenplum安装与部署

第六节 Greenplum使用小技巧

Greenplum介绍

Greenplum 产品发展历程

- Greenplum 公司成立于2003年，产品基于开源的PostgreSQL 数据库研发，2006年推出首款产品
- 2013年Pivotal公司成立后，Greenplum研发团队并入Pivotal研发中心。目前，Greenplum 全球内核研发团队一百多人，遍布美国硅谷，北京，上海以及欧洲，以及 PostgreSQL 数据库社区的核心开发人员

- Greenplum 研发团队将敏捷软件开发方法学引入到分布式数据库的开发中，通过使用站立会议、回顾会议、结对编程、持续集成、测试驱动、单周迭代等敏捷方法建立了高效的快速反馈系统。例如：目前可以实现2个月左右时间高质量合并PostgreSQL内核一个大版本近2000多个commits



2003

GP Release
SQL 99
(2006)

External Table
HA Enhance
SQL 2001
(2007)

WLM
GPCC
SQL 2003
(2008)



In-DB Miner
SAS HPA
Post GIS 1.4
SQL 2008
(2010~2011)



GP Hadoop
Apache HD
Integration
Madlib 2
GPText
(2012~2013)



Orca 优化器
Post GIS 2.0
AO Table
Update
Open Source
(2014~2015)



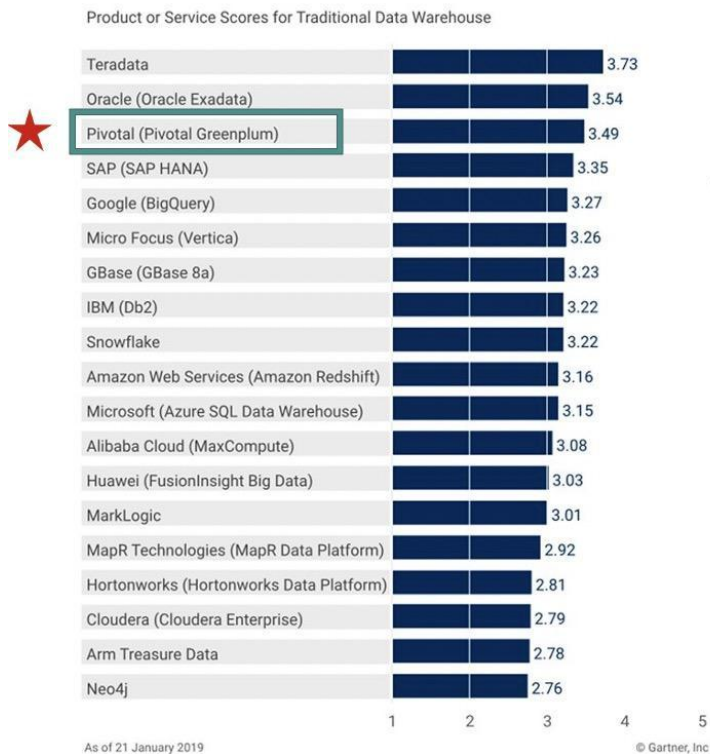
Python, R.....
TensorFlow
Spark, Kafka Integration
HTAP Support
Cloud Support
(2015~至今)



Greenplum市场地位

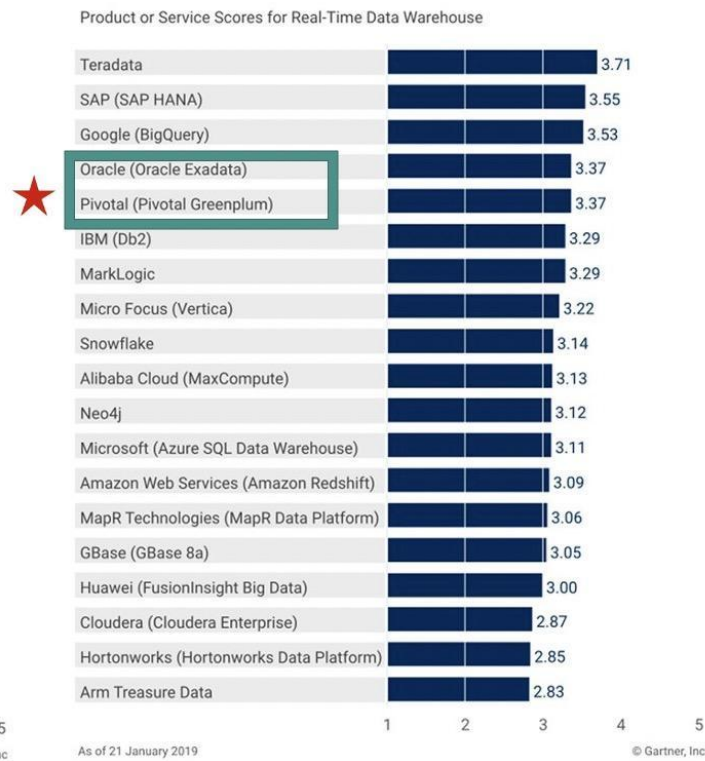
Data Warehouse Use Case(Gartner)

Figure 1. Vendors' Product Scores for Traditional Data Warehouse Use Case



Source: Gartner (March 2019)

Figure 2. Vendors' Product Scores for Real-Time Data Warehouse Use Case



Source: Gartner (March 2019)

- **经典数据分析领域排名第三**

Greenplum在经典数据分析领域排名第三，仅次于Teradata和Oracle。逻辑数据分析领域排名第四

- **实时数据分析领域排名第四**

随着物联网、工业互联网等流式数据分析需求的兴起，实时数据分析能力越来越受重视。Greenplum凭借卓越的性能，在此领域排名和Oracle Exadata并列第四

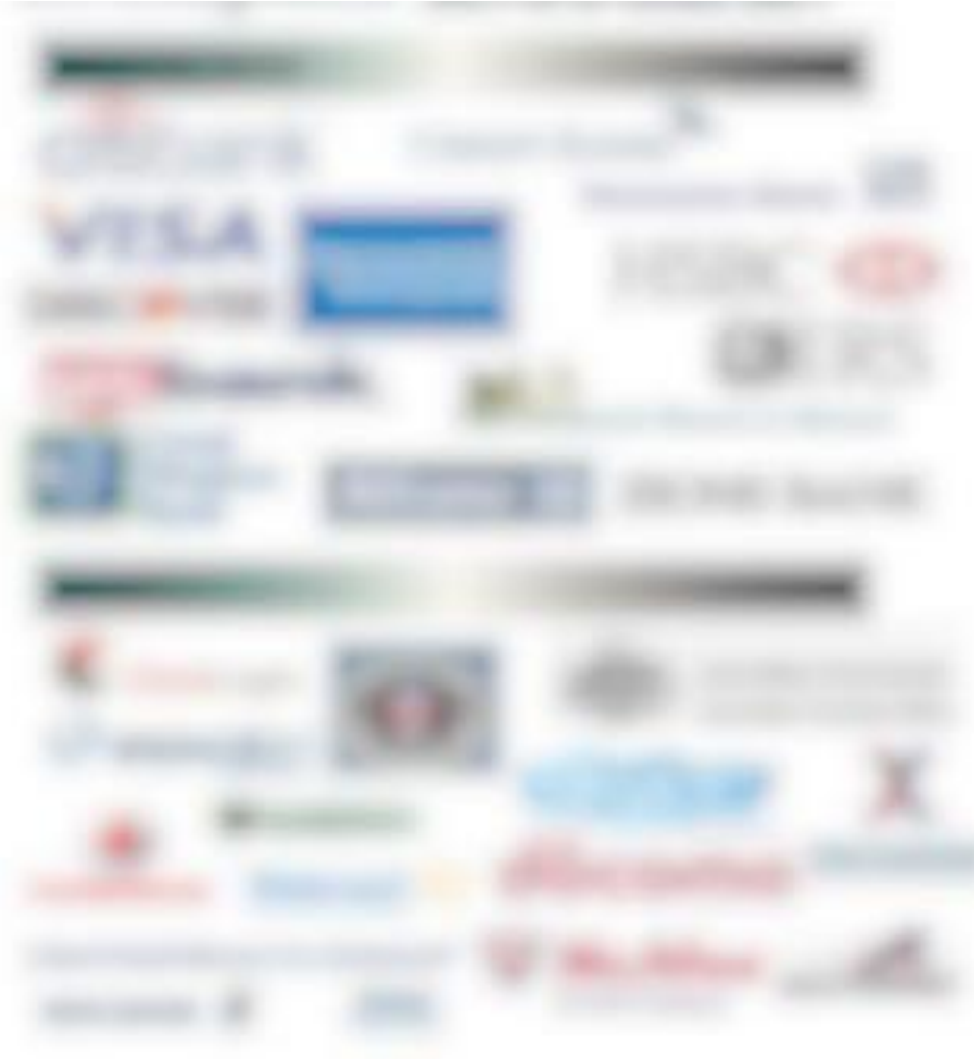
- **前十唯一开源**

Greenplum是全球十大经典和实时数据分析产品中唯一的开源数据库，这就意味着如果选择开源，前十名中别无选择，唯此一款

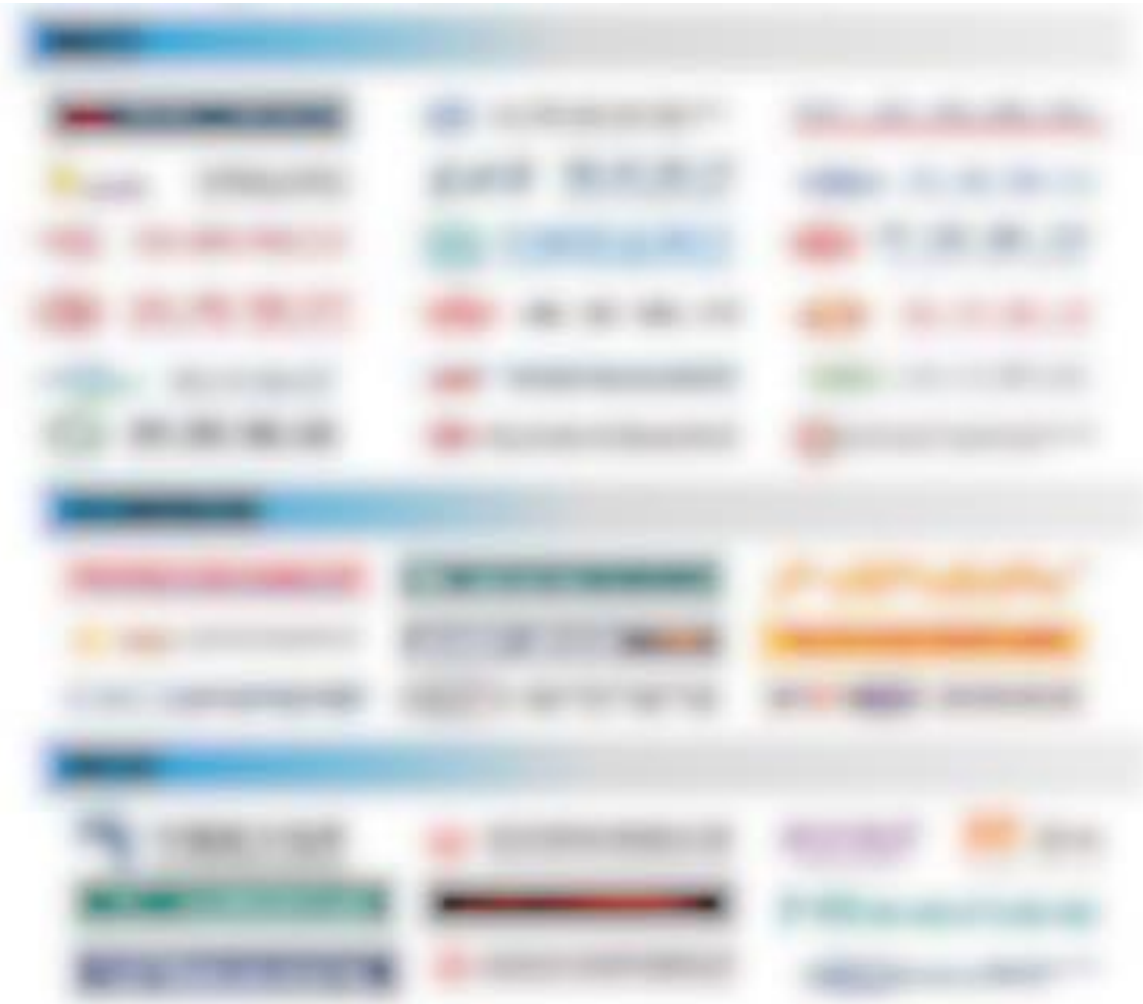
- **“一直被模仿”**

Greenplum是首个商业开源MPP数据库，据中国信息通信研究院数据，参与信通院评测的14款MPP数据库43%都是基于Greenplum

Greenplum 全球典型客户




Greenplum 国内典型客户




Greenplum架构设计

Scale up

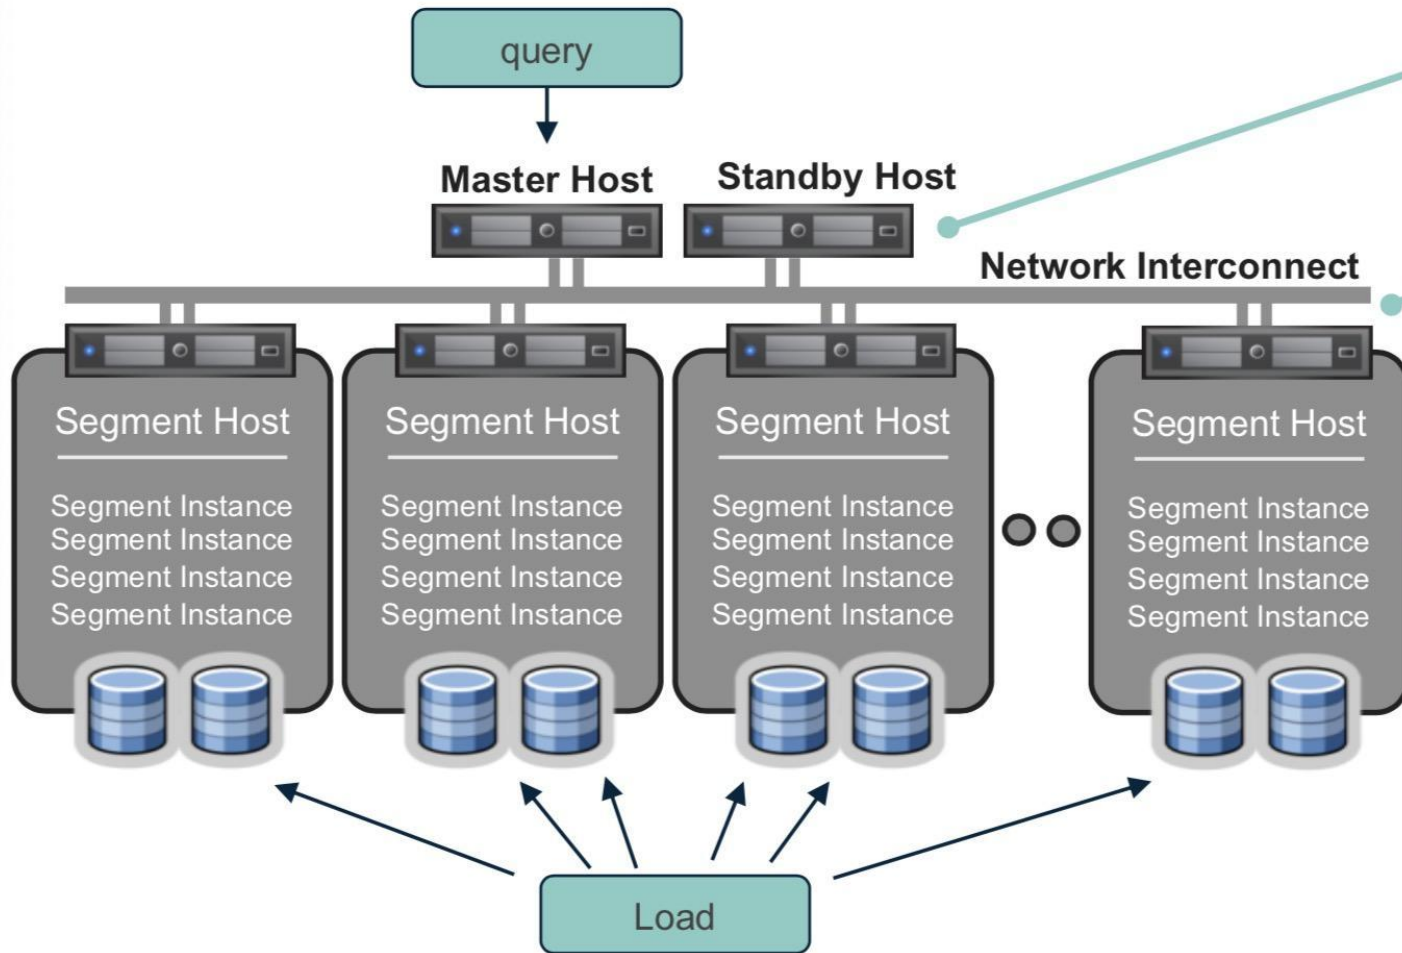


Scale out



Greenplum架构

通过并行处理获得高性能



管理节点与数据节点一起工作

用于连续数据处理流水线的高速互连

- 具有一个或多个实例的数据节点服务器
- 节点上的实例并行处理
- 节点服务器有自己的CPU、磁盘和内存 (Shared Nothing)

Greenplum架构组成



GREENPLUM
DATABASE®

腾讯云大学

Master Hosts

- 系统入口点
- 数据库侦听器进程
- 处理所有用户连接
- 创建查询计划
- 系统管理工具
- 不包含用户数据

Interconnect


- Greenplum数据库连接层
- 元组重新清洗和运输
- 1Gb/10Gb/20Gb网络基础设施
- 私有LAN配置

Segment Hosts

- 每个主机包括用户数据的一部分
- 每个都有自己的CPU、磁盘和内存 (Shared Nothing)
- 用户无法直接访问
- 所有客户端连接都通过Master进入
- 数据库侦听器进程侦听来自服务器的连接

Greenplum DB: 真正完全无共享的MPP数据库

- 真正的完全无共享的并行处理架构, 支持工业标准的X86 服务器
- 数据跨越所有节点均匀分布, 所有节点以并行方式工作, 支持PB级以上的海量存储和处理
- 每个Rack(16节点), 每小时16TB加载性能
- 集群以搭积木方式横向扩展, 目前国内客户单一集群200个节点左右



Greenplum机器选型

机器选型一般配置

序号	产品类型	数量	主要配置要求	
1	计算节点 PC 服务器	X台	外型	2U高机架服务器
			处理器	2颗CPU，每颗12核（启用超线程后每颗24核）、主频不低于2.5G、L3缓存不低于20MB，支持DDR4-1866或以上标准内存。 注：购买时服务器主流配置即可
			内存	256GB RDIMM DDR4-1866或以上标准内存
			硬盘	24块1.2TB、10K PRS、2.5"、热插拔SAS硬盘
			RAID卡	1块RAID卡，不低于双通道，每通道性能不低于6Gb/秒、缓存不低于1GB、支持RAID10和RAID5、支持回写和预读模式、支持电容级掉电保护 注：Raid卡型号建议Megcli OEM
			网络接口	2块非板载、同生产厂商（Intel或博科优先）、同型号万兆光纤网卡，每块包含2个万兆光纤以太网接口，与所配万兆光纤交换机完全兼容
			兼容性	支持RedHat/CentOS 7以上版本
			服务	上架安装服务，3年7×24×4小时生产厂商免费带备件上门维保服务，3年硬盘保留服务
2	万兆光纤 交换机	2台	网络接口	不低于28个万兆以太网光纤接口并满配SFP+模块，与所配PC服务器完全兼容
			服务	上门安装服务，3年7×24×4小时生产厂商免费带备件上门维保服务
3	千兆交换 机	1台	网络接口	不低于48个千兆以太网RJ45接口
			服务	上门安装服务，3年7×24×4小时生产厂商免费带备件上门维保服务

磁盘配置注意事项：

- 留出2块为Hot Spare盘
- 剩下22块盘分为两组并做Raid 5
- 每个RAID组的条带大小都为256KB，写cache策略为“FORCE WRITE BACK”，读磁盘策略设置为“READ AHEAD”

网络配置注意事项：

- 网卡配置建议采用双网卡绑定模式，采用Mode4，支持802.3ad协议，实现动态链路聚合，Active-Active方式，同时需要交换机的链路聚合LACP方式配合支持；
- 千兆交换机仅用于管理；
- 万兆交换机用于集群内部节点通信；

Greenplum安装与部署

系统准备-操作系统



GREENPLUM
DATABASE®

腾讯云大学

https://gpdb.docs.pivotal.io/6-0/install_guide/platform-requirements.html

Operating Systems

Pivotal Greenplum 6 runs on the following operating system platforms:

- Red Hat Enterprise Linux 64-bit 7.x
 - Red Hat Enterprise Linux 64-bit 6.x
 - CentOS 64-bit 7.x
 - CentOS 64-bit 6.x
 - Ubuntu 18.04 LTS
- Resource group on RedHat 6.x and CentOS 6.x: upgrade your kernel to version 2.6.32-696
 - Red Hat Enterprise Linux 7.x or CentOS 7.x prior to 7.3. might cause database hang. RHEL 7.3 and CentOS 7.3 resolves the issue.

Java

Greenplum 6 supports these Java versions for PL/Java and PXF:

- Open JDK 8 or Open JDK 11
- Oracle JDK 8 or Oracle JDK 11

系统准备-最小硬件要求



GREENPLUM
DATABASE®



Table 1. Minimum Hardware Requirements

Minimum CPU	Any x86_64 compatible CPU
Minimum Memory	16 GB RAM per server
Disk Space Requirements	<ul style="list-style-type: none">• 150MB per host for Greenplum installation• Approximately 300MB per segment instance for meta data• Appropriate free space for data with disks at no more than 70% capacity
Network Requirements	<p>10 Gigabit Ethernet within the array</p> <p>NIC bonding is recommended when multiple interfaces are present</p> <p>Pivotal Greenplum can use either IPV4 or IPV6 protocols.</p>

系统准备-存储



GREENPLUM
DATABASE®

腾讯云大学

- GP仅支持XFS 文件系统
- 如果共存储使用块设备存储提供给运行Greenplum数据库的服务器享, 并且挂载到XFS文件系统, 则网络或共享存储支持Greenplum数据库。不支持网络文件系统(NFS)
- Greenplum数据库不直接支持共享存储的其他功能(如重复数据消除或复制), 但只要不干扰Greenplum数据库的预期操作, 就可以在存储供应商的支持下使用这些功能
- Greenplum数据库可以部署在虚拟化系统中, 前提使用块设备存储, 并且可以挂载为XFS文件系统

警告: 在超融合 (HCI) 上运行Greenplum数据库存在性能、可伸缩性和稳定性方面的已知问题, 不建议将其作为关键Greenplum数据库的可伸缩解决方案

容量估算

计算可用磁盘容量

- 磁盘数量 : $\text{disk_size} * \text{number_of_disks}$
- 计算Raid后及格式化后容量 : $(\text{raw_capacity} * 0.9) * \text{number_of_actual_disks}$
- 性能最佳时可用容量 : $\text{formatted_disk_space} * 0.7$
- 配置Mirror及临时空间可用容量 : $(2 * U) + U/3 = \text{usable_disk_space}$
- 压缩比:3:1

$$(24-2)*1.2\text{TB}=26.4\text{TB}$$

$$(24-2-2)*1.2\text{TB}*0.9=21.6\text{TB}$$

$$(24-2-2)*1.2\text{TB}*0.9*0.7= 15.12\text{TB}$$

$$(24-2-2)*1.2\text{TB}*0.9*0.7*3/7\approx 6.5\text{TB}$$

$$6.5\text{T}*3=19.5\text{TB}$$

计算用户数据大小

- Page Overhead
- Row Overhead
- Attribute Overhead
- Indexes

raw data * 1.4

计算元数据和日志大小

- System Metadata
- Write Ahead Log (2 * checkpoint_segments + 1)
- Database Log Files
- Command Center Data

禁用SELinux and Firewall

禁用SELinux

- /etc/selinux/config file. (As root)
SELINUX=disabled

```
# sestatus
SELinuxstatus: disabled
```

禁用防火墙

- Disable firewall software such as iptables (on systems such as RHEL 6.x and CentOS 6.x) or firewalld (on systems such as RHEL 7.x and CentOS 7.x). (As root)

/sbin/chkconfig iptables off

```
# /sbin/chkconfig --list iptables
iptables 0:off 1:off 2:off 3:off 4:off 5:off 6:off
```

- Disable firewalld
systemctl stop firewalld.service
systemctl disable firewalld.service

```
# systemctl status firewalld
* firewalld.service - firewalld - dynamic firewall
daemon Loaded: loaded
(/usr/lib/systemd/system/firewalld.service;
disabled; vendor preset: enabled) Active: inactive
(dead)
```

操作系统内核调优



GREENPLUM
DATABASE®

腾讯云大学

- 编辑/etc/sysctl.conf

```
# kernel.shmall = _PHYS_PAGES / 2 # See Shared Memory Pages
kernel.shmall = 197951838
# kernel.shmmax = kernel.shmall * PAGE_SIZE
kernel.shmmax = 810810728448
kernel.shmni = 4096
vm.overcommit_memory = 2 # See Segment Host Memory
vm.overcommit_ratio = 95 # See Segment Host Memory

net.ipv4.ip_local_port_range = 10000 65535 # See Port Settings
kernel.sem = 500 2048000 200 4096
kernel.sysrq = 1
kernel.core_uses_pid = 1
kernel.msgmnb = 65536
kernel.msgmax = 65536
kernel.msgmni = 2048
net.ipv4.tcp_syncookies = 1
net.ipv4.conf.default.accept_source_route = 0
net.ipv4.tcp_max_syn_backlog = 4096
net.ipv4.conf.all.arp_filter = 1
net.core.netdev_max_backlog = 10000
net.core.rmem_max = 2097152
net.core.wmem_max = 2097152
vm.swappiness = 10
vm.zone_reclaim_mode = 0
vm.dirty_expire_centisecs = 500
vm.dirty_writeback_centisecs = 100
vm.dirty_background_ratio = 0 # See System Memory
vm.dirty_ratio = 0
vm.dirty_background_bytes = 1610612736
vm.dirty_bytes = 4294967296
```

Shared_buffers > 16K * *max_connections*

SHMMAX=shared_buffers + other_seg_shmem

SHMALL=(num_instances_per_host * (shared_buffers + other_seg_shmem)) + other_app_shared_mem

SHMMAX

echo \$(expr \$(getconf _PHYS_PAGES) / 2)

SHMALL

echo \$(expr \$(getconf _PHYS_PAGES) / 2 * \$(getconf PAGE_SIZE))

操作系统内核调优



GREENPLUM
DATABASE®

腾讯云大学

- 编辑/etc/sysctl.conf

```
# kernel.shmall = _PHYS_PAGES / 2 # See Shared Memory Pages
kernel.shmall = 197951838
# kernel.shmmax = kernel.shmall * PAGE_SIZE
kernel.shmmax = 810810728448
kernel.shmni = 4096
vm.overcommit_memory = 2 # See Segment Host Memory
vm.overcommit_ratio = 95 # See Segment Host Memory

net.ipv4.ip_local_port_range = 10000 65535 # See Port Settings
kernel.sem = 500 2048000 200 4096
kernel.sysrq = 1
kernel.core_uses_pid = 1
kernel.msgmnb = 65536
kernel.msgmax = 65536
kernel.msgmni = 2048
net.ipv4.tcp_syncookies = 1
net.ipv4.conf.default.accept_source_route = 0
net.ipv4.tcp_max_syn_backlog = 4096
net.ipv4.conf.all.arp_filter = 1
net.core.netdev_max_backlog = 10000
net.core.rmem_max = 2097152
net.core.wmem_max = 2097152
vm.swappiness = 10
vm.zone_reclaim_mode = 0
vm.dirty_expire_centisecs = 500
vm.dirty_writeback_centisecs = 100
vm.dirty_background_ratio = 0 # See System Memory
vm.dirty_ratio = 0
vm.dirty_background_bytes = 1610612736
vm.dirty_bytes = 4294967296
```

vm.overcommit_memory = 2
vm.overcommit_ratio = 95

示例: RAM 128GB ,Swap 128GB

Memory Allocation Limit = 128 GB Swap Space + 128 GB RAM * (95
Overcommit Ratio / 100)

Memory Allocation Limit = 249.6 GB

- More than 64GB

vm.dirty_background_ratio = 0

vm.dirty_ratio = 0

vm.dirty_background_bytes = 1610612736 # 1.5GB

vm.dirty_bytes = 4294967296 # 4GB

- With 64GB of memory or less

vm.dirty_background_ratio = 3

vm.dirty_ratio = 10

时钟设置



GREENPLUM
DATABASE®

腾讯云大学

配置NTP

- On the master host:
server XX.XX.XX.XX
- On each segment host:
server mdw prefer
server smdw
- On the standby master host:
server mdw prefer
server XX.XX.XX.XX
- 验证NTP
gpssh -f hostfile_gpssh_allhosts -v -e 'ntpd'

系统资源限制



GREENPLUM
DATABASE®

腾讯云大学

- 修改 /etc/security/limits.conf
 - * soft nofile 1048576
 - * hard nofile 1048576
 - * soft nproc 1048576
 - * hard nproc 1048576
- 修改/etc/security/limits.d/90-nproc.conf file (RHEL/CentOS 6)
/etc/security/limits.d/20-nproc.conf file (RHEL/CentOS 7)
 - * soft nproc 1048576
 - * hard nproc 1048576
 - * soft nofile 1048576
 - * hard nofile 1048576

磁盘 I/O 及其它参数



GREENPLUM
DATABASE®

腾讯云大学

- 挂载 XFS 文件系统
rw,nodev,noatime,nobarrier,inode64
- 设置read-ahead
`/sbin/blockdev --setra 16384 devname`
- 设置I/O 调度策略
`echo deadline > /sys/block/devname/queue/scheduler`
`grubby --update-kernel=ALL --args="elevator=deadline"`
- 禁用 Transparent Huge Pages (THP)
`grubby --update-kernel=ALL --args="transparent_hugepage=never"`
- 设置RemoveIPC
`/etc/systemd/logind.conf`
RemoveIPC=no
- 设置SSH连接阈值
Max Startups 10000:30:20000

创建用户



GREENPLUM
DATABASE®



腾讯云大学

- 创建组

```
groupadd -g 599 gpadmin
```

- 创建用户

```
useradd -g gpadmin -u 600 gpadmin  
echo "password" | passwd --stdin
```


Greenplum软件安装



GREENPLUM
DATABASE®



腾讯云大学

- 商业版

<https://network.pivotal.io/products/pivotal-gpdb/>

yum install ./greenplum-db-<version>-<platform>.rpm

- 开源版

<https://github.com/greenplum-db/gpdb>

<https://github.com/greenplum-db/gpdb/blob/master/README.linux.md>

./configure --with-perl --with-python --with-libxml --with-gssapi --prefix=/usr/local/gpdb

make -j8

make -j8 install

```
/bin/mkdir -p '/usr/local/gpdb/lib/postgresql'
ln -sf gpcloud.so /usr/local/gpdb/lib/postgresql/gps3ext.so
/bin/install -c -m 755 gpcloud.so '/usr/local/gpdb/lib/postgresql/gpcloud.so'
make[2]: Leaving directory '/home/gpadmin/gpdb-master/gpcontrib/gpcloud'
make[2]: Entering directory '/home/gpadmin/gpdb-master/gpcontrib/gpcloud/bin/gpcheckcloud'
/bin/mkdir -p '/usr/local/gpdb/bin'
/bin/install -c gpcheckcloud '/usr/local/gpdb/bin'
make[2]: Leaving directory '/home/gpadmin/gpdb-master/gpcontrib/gpcloud/bin/gpcheckcloud'
make[2]: Entering directory '/home/gpadmin/gpdb-master/gpcontrib/pxf_fdw'
/bin/mkdir -p '/usr/local/gpdb/lib/postgresql'
/bin/mkdir -p '/usr/local/gpdb/share/postgresql/extension'
/bin/mkdir -p '/usr/local/gpdb/share/postgresql/extension'
/bin/install -c -m 755 pxf_fdw.so '/usr/local/gpdb/lib/postgresql/pxf_fdw.so'
/bin/install -c -m 644 ./pxf_fdw.control '/usr/local/gpdb/share/postgresql/extension/'
/bin/install -c -m 644 ./pxf_fdw--1.0.sql '/usr/local/gpdb/share/postgresql/extension/'
make[2]: Leaving directory '/home/gpadmin/gpdb-master/gpcontrib/pxf_fdw'
make[1]: Leaving directory '/home/gpadmin/gpdb-master/gpcontrib'
Greenplum Database installation complete.
```

```
[gpadmin@gpthird ~]$ ls -l /usr/local/gpdb/
total 16
drwxr-xr-x 7 gpadmin gpadmin 4096 Jun 29 07:54 bin
drwxr-xr-x 3 gpadmin gpadmin  22 Jun 29 07:54 docs
-rw-r--r-- 1 gpadmin gpadmin  336 Jun 29 07:54 greenplum_path.sh
drwxr-xr-x 4 gpadmin gpadmin 4096 Jun 29 07:54 include
drwxr-xr-x 5 gpadmin gpadmin 4096 Jun 29 07:54 lib
drwxr-xr-x 2 gpadmin gpadmin  21 Jun 29 07:54 libexec
drwxr-xr-x 2 gpadmin gpadmin  176 Jun 29 07:54 sbin
drwxr-xr-x 4 gpadmin gpadmin   41 Jun 29 07:54 share
```

Greenplum软件安装(建立互信和目录)



GREENPLUM
DATABASE®

腾讯云大学

- 确认GP软件安装成功并使用gpadmin用户登录
source /usr/local/greenplum-db/greenplum_path.sh
- 确认所有服务器/etc/hosts包含各主机名, 并创建一个包含所有主机名的文件all_hosts

```
mdw  
smdw  
sdw1  
sdw2  
sdw3  
sdw4
```

- 使用gpssh-exkeys 工具建立互信
gpssh-exkeys -f all_hosts
- 使用gpssh 工具登录无输入密码提示
gpssh -f all_hosts -e '-ls \$GPHOME'
- 在master & standby master 创建数据目录
mkdir /data/master
chown gpadmin /data/master
- 在所有segment主机创建数据目录
gpssh -f all_segs -e 'mkdir /data/primary'
gpssh -f all_segs -e 'mkdir /data/mirror'
gpssh -f all_segs -e 'chown gpadmin /data/primary'
gpssh -f all_segs -e 'chown gpadmin /data/mirror'

Greenplum软件安装(检验性能)



GREENPLUM
DATABASE®

腾讯云大学

- 检验Disk I/O 性能和 内存带宽

```
gpcheckperf -f hostfile_gpcheckperf -r ds -D -d /data/primary -d /data/mirror
```

- 检验网络性能

```
gpcheckperf -f hostfile_gpchecknet_ic -r N -d /tmp > subnet.out
```

```
gpcheckperf -f hostfile_gpchecknet_ic -r M --duration=3m -d /tmp > checknet.m.log
```

Greenplum数据库初始化

- 创建数据库初始化文件

```
cp $GPHOME/docs/cli_help/gpconfigs/gpinitsystem_config  
~/gpconfigs/gpinitsystem_config  
然后编辑~/gpconfigs/gpinitsystem_config
```

```
ARRAY_NAME="EMC Greenplum DW"  
SEG_PREFIX=gpseg PORT_BASE=40000  
declare -a DATA_DIRECTORY=(/data/primary  
/data/primary /data/primary)  
MASTER_HOSTNAME=mdw  
MASTER_DIRECTORY=/data/master  
MASTER_PORT=5432  
TRUSTED_SHELL=ssh  
CHECK_POINT_SEGMENTS=8  
ENCODING=UNICODE
```

```
### Optional (for mirror)  
MIRROR_PORT_BASE=7000  
REPLICATION_PORT_BASE=8000  
MIRROR_REPLICATION_PORT_BASE=9000  
declare -a MIRROR_DATA_DIRECTORY=(/data/mirror  
/data/mirror /data/mirror)
```

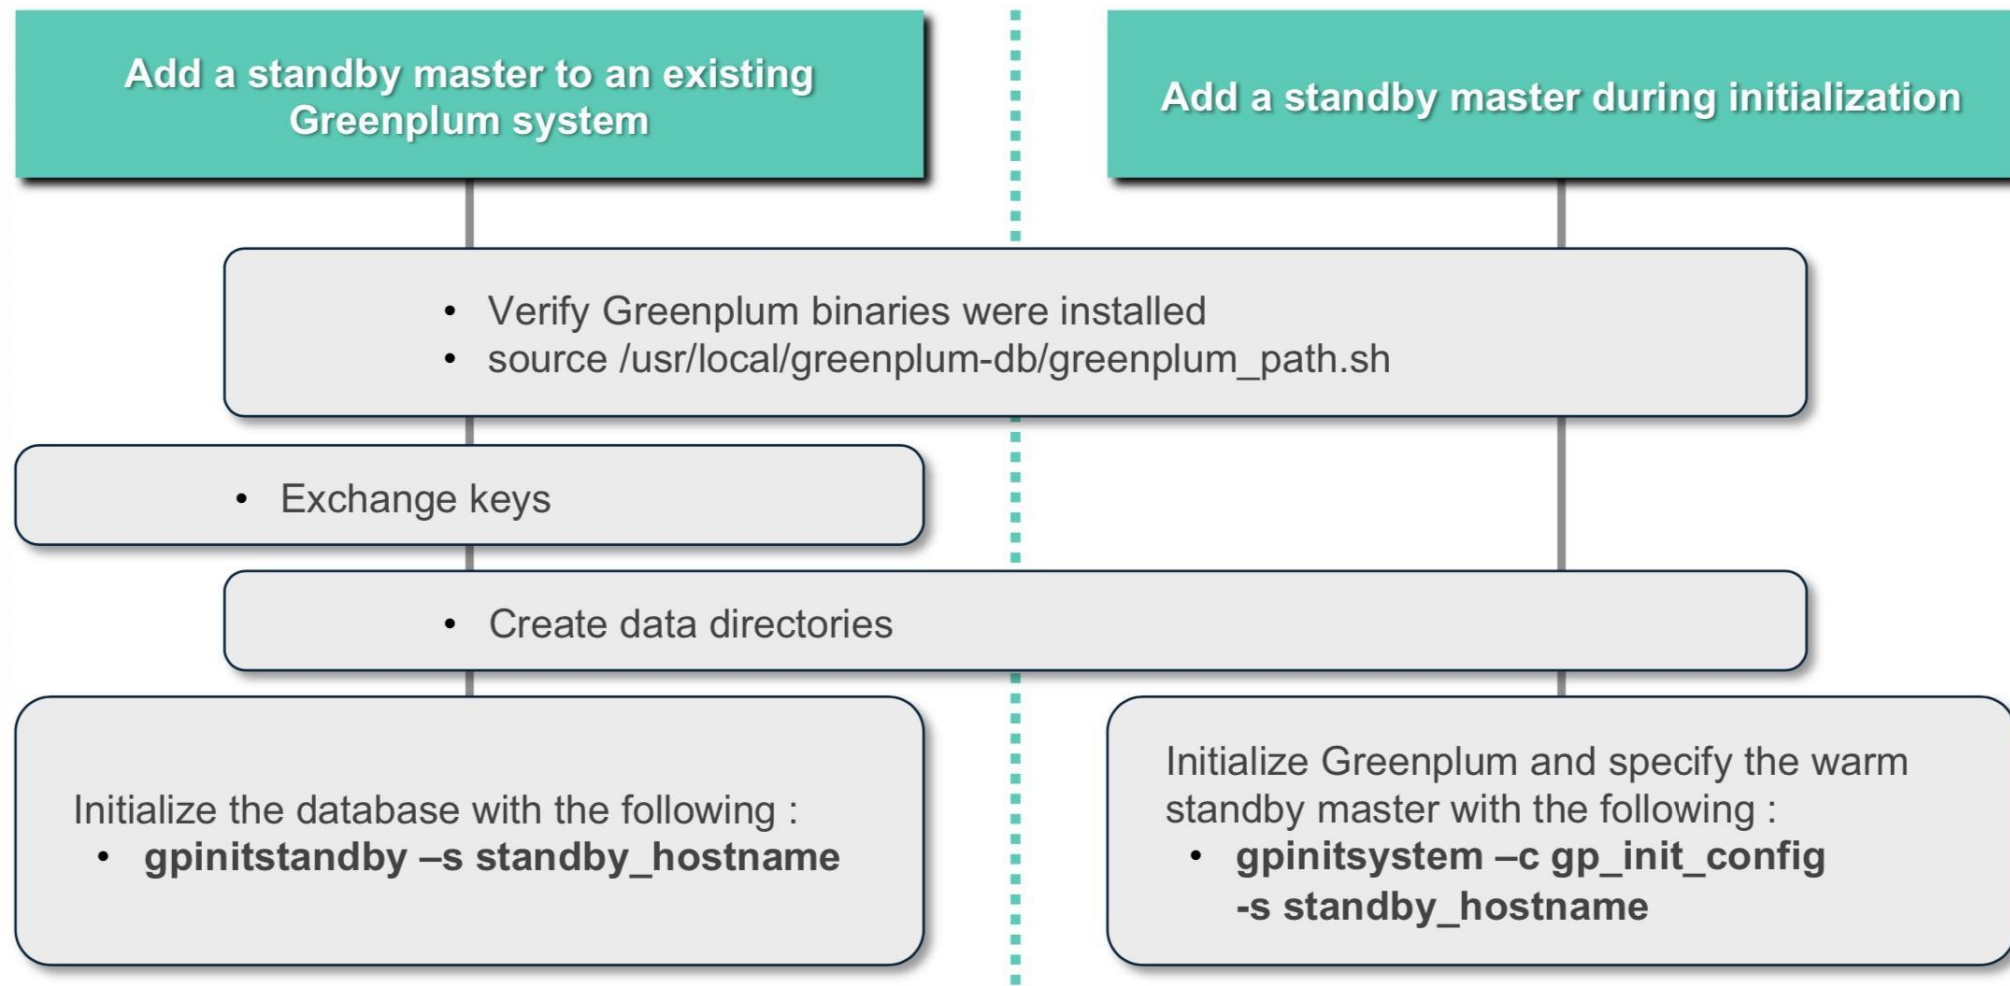
```
declare -a DATA_DIRECTORY =
```

- 运行初始化命令

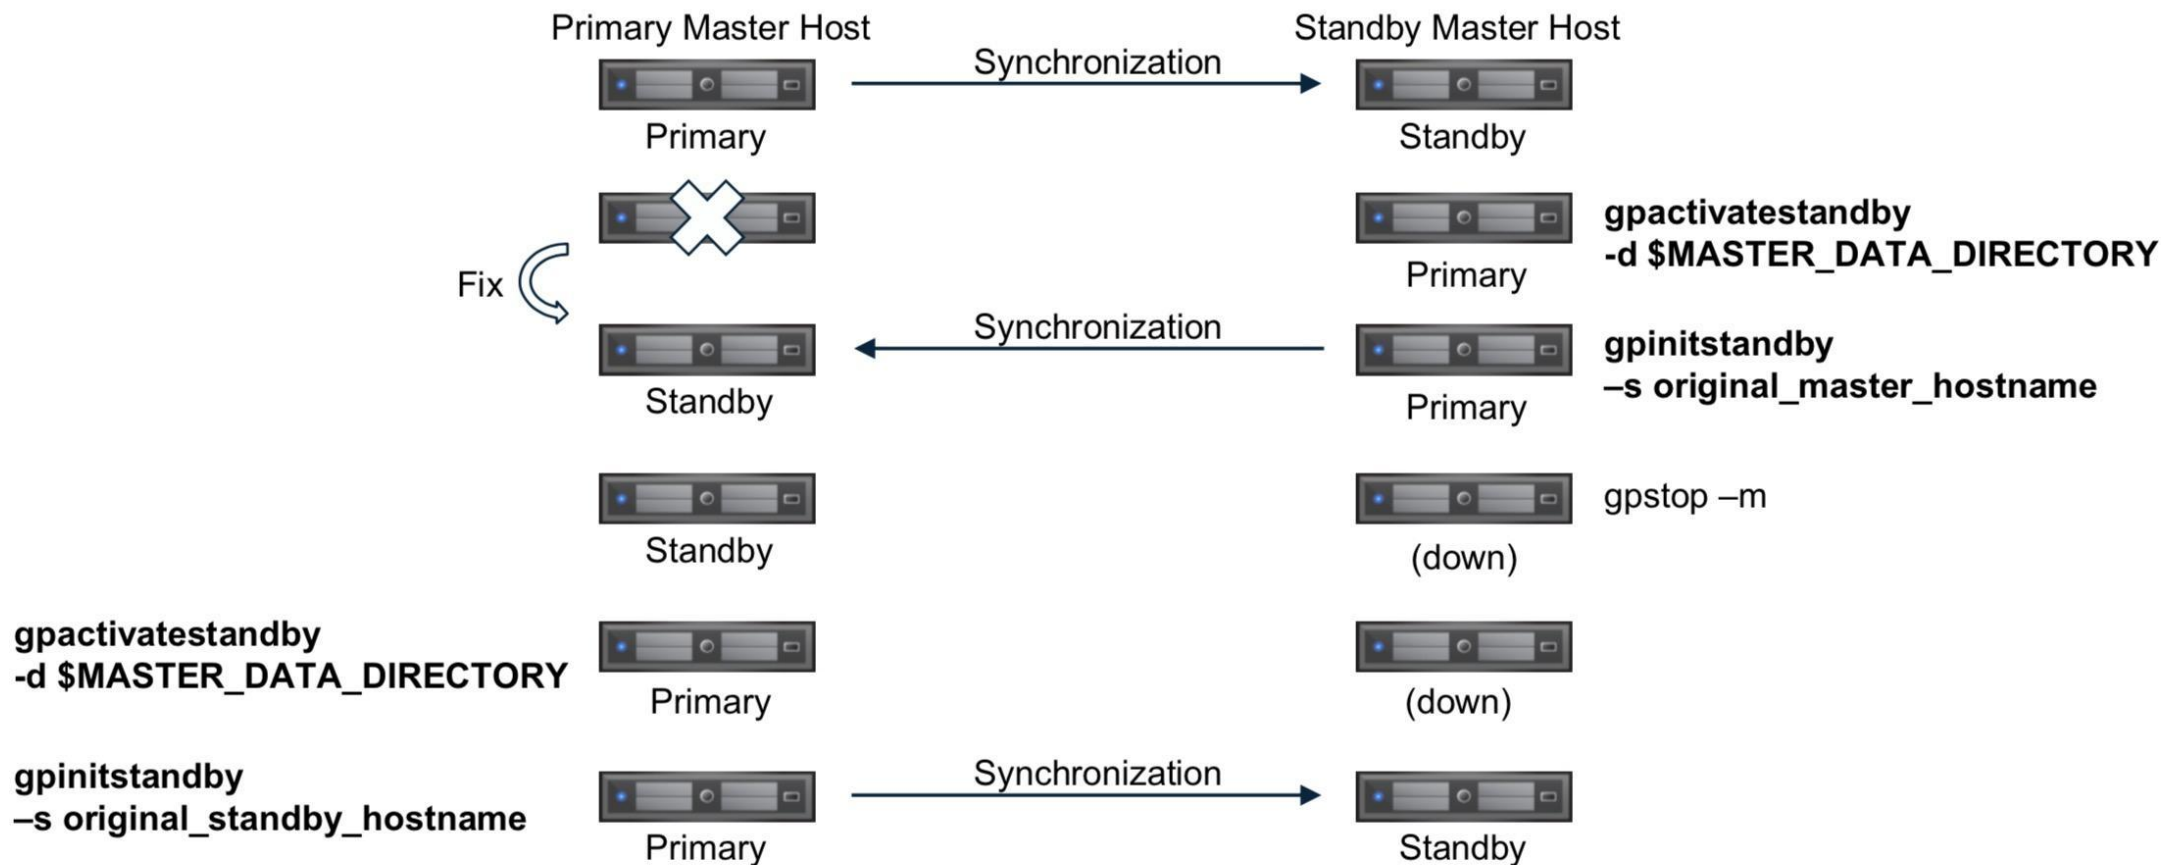
```
gpinitsystem -c gpconfigs/gpinitsystem_config -h gpconfigs/hostfile_gpinitsystem  
或者
```

```
gpinitsystem -c gpconfigs/gpinitsystem_config -h gpconfigs/hostfile_gpinitsystem -s  
standby_master_hostname
```

配置Standby



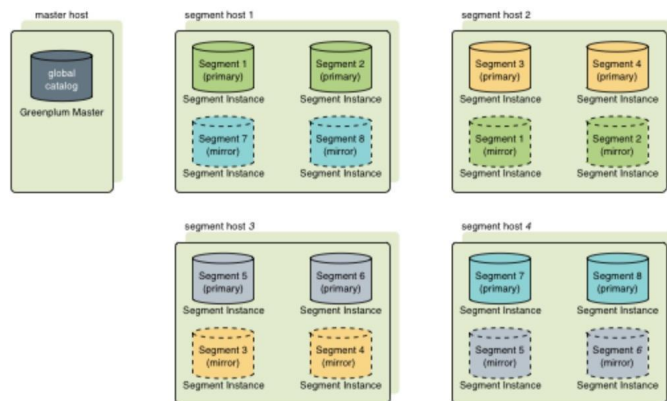
Master Failover和Restoration



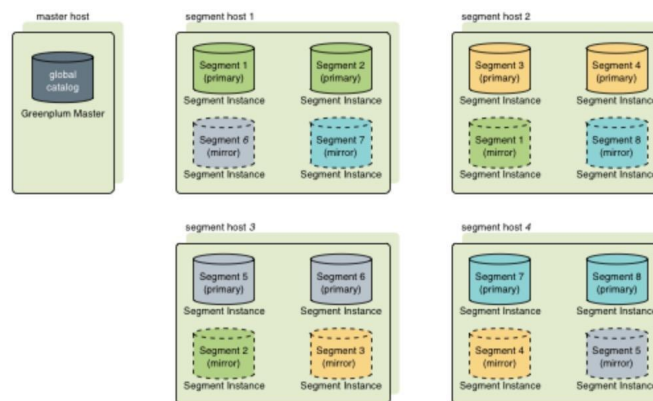
配置Segment主实例镜像

- 镜像是主实例的副本，用于高可用；
- 初始化数据库时可以启用
- 亦可使用gpaddmirrors -i config_file(gpaddmirrors -o)
- 镜像分布策略

1)group



2)spread



3)自定义

配置环境变量

- Master数据目录
MASTER_DATA_DIRECTORY=/data/master/gpseg-1
- GP基础目录
GPHOME=/usr/local/greenplum-db
- 默认登录数据库名
PGDATABASE=edw
- 默认登录端口
PGPORT=5432
- GP环境变量
source the **/usr/local/greenplum-db/greenplum_path.sh**

Greenplum使用小技巧

性能测试参考值



IO读

```
disk read avg time (sec): 77.66
disk read tot bytes: 4322078228480
disk read tot bandwidth (MB/s): 53165.65
disk read min bandwidth (MB/s): 3181.62 [dwtestdn010]
disk read max bandwidth (MB/s): 3561.67 [dwtestdn012]
-- per host bandwidth --
disk read bandwidth (MB/s): 3207.77 [dwtestdn014]
disk read bandwidth (MB/s): 3309.98 [dwtestdn003]
disk read bandwidth (MB/s): 3205.77 [dwtestdn011]
disk read bandwidth (MB/s): 3208.57 [dwtestdn005]
disk read bandwidth (MB/s): 3509.28 [dwtestdn008]
disk read bandwidth (MB/s): 3216.58 [dwtestdn009]
disk read bandwidth (MB/s): 3520.79 [dwtestdn006]
disk read bandwidth (MB/s): 3540.63 [dwtestdn001]
disk read bandwidth (MB/s): 3329.23 [dwtestdn013]
disk read bandwidth (MB/s): 3561.67 [dwtestdn012]
disk read bandwidth (MB/s): 3285.50 [dwtestdn002]
disk read bandwidth (MB/s): 3181.62 [dwtestdn010]
disk read bandwidth (MB/s): 3193.85 [dwtestdn004]
disk read bandwidth (MB/s): 3229.08 [dwtestdn016]
disk read bandwidth (MB/s): 3196.23 [dwtestdn015]
disk read bandwidth (MB/s): 3469.11 [dwtestdn007]
```

IO写

```
disk write avg time (sec): 100.01
disk write tot bytes: 4322078228480
disk write tot bandwidth (MB/s): 41390.30
disk write min bandwidth (MB/s): 2426.90 [dwtestdn005]
disk write max bandwidth (MB/s): 2808.42 [dwtestdn002]
-- per host bandwidth --
disk write bandwidth (MB/s): 2718.90 [dwtestdn003]
disk write bandwidth (MB/s): 2808.42 [dwtestdn002]
disk write bandwidth (MB/s): 2806.89 [dwtestdn007]
disk write bandwidth (MB/s): 2426.90 [dwtestdn005]
disk write bandwidth (MB/s): 2805.05 [dwtestdn008]
disk write bandwidth (MB/s): 2442.32 [dwtestdn009]
disk write bandwidth (MB/s): 2805.97 [dwtestdn006]
disk write bandwidth (MB/s): 2805.05 [dwtestdn001]
disk write bandwidth (MB/s): 2439.08 [dwtestdn013]
disk write bandwidth (MB/s): 2446.26 [dwtestdn012]
disk write bandwidth (MB/s): 2438.16 [dwtestdn011]
disk write bandwidth (MB/s): 2441.16 [dwtestdn010]
disk write bandwidth (MB/s): 2693.04 [dwtestdn004]
disk write bandwidth (MB/s): 2438.39 [dwtestdn016]
disk write bandwidth (MB/s): 2435.16 [dwtestdn015]
disk write bandwidth (MB/s): 2439.54 [dwtestdn014]
```

网络接收发送

```
Per host transfer rates
dwtestdn014 Tx rate: 2128.10
dwtestdn010 Tx rate: 2117.70
dwtestdn011 Tx rate: 2233.52
dwtestdn005 Tx rate: 2017.41
dwtestdn008 Tx rate: 2237.13
dwtestdn009 Tx rate: 2129.66
dwtestdn001 Tx rate: 2237.49
dwtestdn015 Tx rate: 2258.47
.....
dwtestdn016 Tx rate: 1885.08
dwtestdn006 Tx rate: 1741.50
dwtestdn007 Tx rate: 2188.07

Per host receive rates
dwtestdn003 Rx rate: 2030.89
dwtestdn002 Rx rate: 2055.64
dwtestdn016 Rx rate: 2111.35
dwtestdn008 Rx rate: 2199.26
.....
dwtestdn010 Rx rate: 2101.51
dwtestdn004 Rx rate: 2244.87
dwtestdn005 Rx rate: 2232.87
dwtestdn015 Rx rate: 2082.19
dwtestdn014 Rx rate: 2124.66
```

- 硬件问题早发现早解决
- 安装完毕后重启所有服务器

日志输出与查看

- -D --debug或 -v -- verbose 详细日志输出
- GP命令中利用print函数打印变量值

- \$MASTER_DATA_DIRECTORY/pg_log/startup.log
- \$MASTER_DATA_DIRECTORY/pg_log/*.csv

程序调试、监控工具安装



GREENPLUM
DATABASE®



腾讯云大学

- strace
- pstat
- gcore
- gdb
- nmon
- netperf
- netserver

- packcore
- gpmt

- gpcc

常见问题

- 切换用户环境变量

```
su -
```

- RH 6/CentOS 6防火墙禁掉后服务器重启后又Active

```
chkconfig libvirtd off
```

- 磁盘读写性能

```
vm.dirty_background_bytes = 1610612736 # 1.5GB
```

```
vm.dirty_bytes = 4294967296 # 4GB
```

感谢观看



Greenplum中文社区公众号



Greenplum微信技术讨论群