


공개SW 솔루션 설치 & 활용 가이드

시스템SW > 스토리지



제대로 배워보자

How to Use Open Source Software

Open Source Software Installation & Application Guide




오픈소스 소프트웨어 통합지원센터
Open Source Software Support Center



CONTENTS

1. 개요
2. 기능요약
3. 실행환경
4. 설치 및 실행
5. 기능소개
6. . 활용예제
7. . FAQ
8. . 용어정리


1. 개요



| | | | |
|------------------|--|-----------------|---|
| 소개 | <ul style="list-style-type: none"> • Gluster FS는 확장성이 가능한 NAS 파일 시스템 • 다수의 스토리지를 이더넷(Ethernet) 또는 인피니밴드(InfiniBand)를 통하여 하나의 커다란 병렬 네트워크 파일 시스템으로 통합 | | |
| 주요기능 | <ul style="list-style-type: none"> • 로컬 파일 시스템 위에 분산 파일 시스템 구성 • 탄력적인 해싱 알고리즘으로 메타데이터 서버가 필요 없음 • brick 이라는 XFS파일시스템의 디렉터리 단위로 파일이 분산 저장 | | |
| 대분류 | <ul style="list-style-type: none"> • 시스템 SW | 소분류 | <ul style="list-style-type: none"> • 스토리지 |
| 라이선스 형태 | <ul style="list-style-type: none"> • GPL V3 | 사전설치 솔루션 | <ul style="list-style-type: none"> • Linux (O/S) |
| 실행 하드웨어 | <ul style="list-style-type: none"> • 2소켓(4코어,6코어또는8코어)권장 • O/S 설치 디스크 용량 200GB • RAID 6 및 RAID 1 + 0 하드웨어 RAID 컨트롤러 지원 | 버전 | <ul style="list-style-type: none"> • glusterfs-4.1.5-1(2018년 10월 기준) |
| 특징 | <ul style="list-style-type: none"> • 고성능 및 고가용 파일 및 오브젝트 기반 스토리지 구축 • 데이터의 저장 및 접근을 즉시 파일 혹은 오브젝트 기반으로 제공 • 메타데이터 서버를 제거하여 성능, 병목 및 단일장애지점 방지 | | |
| 보안취약점 | <ul style="list-style-type: none"> • 취약점 ID : CVE-2018-1000199 • 심각도 : 5.5 MEDIUM(V3) • 취약점 설명 : Linux Kernel 버전 3.18에는 크래시 및 메모리 손상을 초래할 수 있는 modify_user_hw_breakpoint ()에 위험한 기능의 취약점 발견 • 대응방안 : 최신 패치 적용 • 참고 경로 : https://www.securitytracker.com/id/1040806 | | |
| 개발회사/커뮤니티 | <ul style="list-style-type: none"> • https://www.gluster.org/community/ | | |
| 공식 홈페이지 | <ul style="list-style-type: none"> • https://www.gluster.org/ | | |



2. 기능요약




- Gluster fs 기능

| 기능 | 지원여부 |
|---------------------------------|------|
| distributed | 지원 |
| replicated | 지원 |
| striped | 지원 |
| distributed striped | 지원 |
| distributed replicated | 지원 |
| distributed striped replication | 지원 |
| striped replicated | 지원 |



3. 실행환경



- Use Case에 따라 하드웨어 사양이 달라짐

| 구분 | HPC | General Purpose | Archival |
|---------------|---|---|---|
| Cpu(이상) | 2-socket (with 4-core, 6-core, or 8-core) | 2-socket (with 4-core, 6-core, or 8-core) | 2-socket (with 4-core, 6-core, or 8-core) |
| Memory(최소) | 48GB | 32GB | 16GB |
| 네트워크 | 2x10GigE | 2x10GigE (preferred) or 2x1GigE | 2x10GigE (preferred) or 2x1GigE |
| Hard Disk 사양 | 15000 RPM drives(2.5" inch SAS) OR Solid state disks. | 7200 or 10000 RPM drives (3.5" SAS or SATA) | 7200 or 10000 RPM drives (3.5" SAS or SATA) |
| 파일 시스템 크기(최소) | 900GB | 8 TB | 8 TB |

※ 위 내용은 권장 사항이 아니며, workload에 따라 하드웨어 사양이 변경 될 수 있음



4. 설치 및 실행




세부 목차

1. 설치 환경 준비
2. 설치 진행
3. 설치 완료



4. 설치 및 실행




4.1 설치 환경 준비

- Gluster fs 설치 환경

| 하드웨어 및 소프트웨어 정보 | 비고 |
|------------------------------|-----------------------------|
| Gluster fs 서버 | 최소 2대(구성환경에 따라 서버 대수가 달라짐) |
| O/S 정보 | Linux |
| O/S 설치 공간 | 200GB 권장(구성환경에 따라 달라짐) |
| Gluster fs 구성에 필요한 파일 시스템 영역 | 권장은 없으며, 구성환경에 따라 달라짐 |
| 네트워크 | 네트워크 통신에 있어서 최대한의 대역폭 확보 필요 |



4. 설치 및 실행



4.2 설치 진행(1/5)

- **Gluster fs를 구성하기 위해 리눅스(O/S) 설치 및 네트워크 설정 진행(2 EA)**
 - O/S 버전 : Cent OS 7.5
 - hostname : gluster01, gluster02
 - 네트워크 설정 : gluster01 : 10.65.10.101
gluster02 : 10.65.10.102

```
root@gluster01:~# hostname
gluster01
root@gluster01:~# cat /etc/redhat-release
CentOS Linux release 7.5.1804 (Core)
root@gluster01:~# ifconfig eth0
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 10.65.10.101 netmask 255.255.255.0 broadcast 10.65.10.255
    inet6 fe80::5054:ff:fed1:71e prefixlen 64 scopeid 0x20<link>
    ether 52:54:00:d1:07:1e txqueuelen 1000 (Ethernet)
    RX packets 107733 bytes 23420299 (22.3 MiB)
    RX errors 0 dropped 55700 overruns 0 frame 0
    TX packets 5742 bytes 581802 (568.1 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

root@gluster01:~#

root@gluster02:~# hostname
gluster02
root@gluster02:~# cat /etc/redhat-release
CentOS Linux release 7.5.1804 (Core)
root@gluster02:~# ifconfig eth0
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 10.65.10.102 netmask 255.255.255.0 broadcast 10.65.10.255
    inet6 fe80::5054:ff:feeb:49ad prefixlen 64 scopeid 0x20<link>
    ether 52:54:00:eb:49:ad txqueuelen 1000 (Ethernet)
    RX packets 150132 bytes 327435583 (312.2 MiB)
    RX errors 0 dropped 55702 overruns 0 frame 0
    TX packets 43190 bytes 2909990 (2.7 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

root@gluster02:~#
```

- **/etc/hosts 파일 설정 및 ntp 서버에 시간 동기화 설정**
 - /etc/hosts 파일 설정 : 10.65.10.101 gluster01 => 추가
10.65.10.102 gluster02 => 추가
 - ntp 서버 시간 동기화 : time.bora.net



4. 설치 및 실행



4.2 설치 진행(2/5)


```
root@gluster01~  
[root@gluster01 ~]#  
[root@gluster01 ~]#  
[root@gluster01 ~]# cat /etc/hosts  
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4  
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6  
10.65.10.101 gluster01  
10.65.10.102 gluster02  
[root@gluster01 ~]# ntpq -p  
remote refid st t when poll reach delay offset jitter  
-----  
*time.bora.net 204.123.2.5 2 u 37 64 1 2.975 1.473 1.327  
[root@gluster01 ~]# date  
Tue Oct 16 16:07:01 KST 2018  
[root@gluster01 ~]#  
  
root@gluster02~  
[root@gluster02 ~]# cat /etc/hosts  
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4  
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6  
10.65.10.101 gluster01  
10.65.10.102 gluster02  
[root@gluster02 ~]# ntpq -p  
remote refid st t when poll reach delay offset jitter  
-----  
*time.bora.net 204.123.2.5 2 u 3 64 1 4.190 1.611 1.867  
[root@gluster02 ~]# date  
Tue Oct 16 16:07:13 KST 2018  
[root@gluster02 ~]#
```

- **Gluster fs 구성에 필요한 파일 시스템 영역 구성**

- gluster01 서버 : 10GB 볼륨 마운트 : 2EA => /dev/vdb1 /gluster01/data1
/dev/vdc1 /gluster01/data2
- gluster02 서버 : 10GB 볼륨 마운트 : 2EA => /dev/vdb1 /gluster02/data1
/dev/vdc1 /gluster02/data2



4. 설치 및 실행



4.2 설치 진행(3/5)


```
root@gluster01:gluster01
[root@gluster01 gluster01]# hostname
gluster01
[root@gluster01 gluster01]# df -Th
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda3       xfs       21G   3.5G  18G  17% /
devtmpfs        devtmpfs  2.0G   0     2.0G   0% /dev
tmpfs           tmpfs     2.0G   0     2.0G   0% /dev/shm
tmpfs           tmpfs     2.0G   9.2M   2.0G   1% /run
tmpfs           tmpfs     2.0G   0     2.0G   0% /sys/fs/cgroup
/dev/vda1       xfs      1014M  155M  860M  16% /boot
tmpfs           tmpfs     396M   12K  396M   1% /run/user/42
tmpfs           tmpfs     396M   0     396M   0% /run/user/0
/dev/vdb1       xfs       10G   33M   10G   1% /gluster01/data1
/dev/vdc1       xfs       10G   33M   10G   1% /gluster01/data2
[root@gluster01 gluster01]#

root@gluster02:gluster02
[root@gluster02 gluster02]# hostname
gluster02
[root@gluster02 gluster02]# df -Th
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda3       xfs       21G   3.8G  18G  18% /
devtmpfs        devtmpfs  2.0G   0     2.0G   0% /dev
tmpfs           tmpfs     2.0G   0     2.0G   0% /dev/shm
tmpfs           tmpfs     2.0G   9.4M   2.0G   1% /run
tmpfs           tmpfs     2.0G   0     2.0G   0% /sys/fs/cgroup
/dev/vda1       xfs      1014M  155M  860M  16% /boot
tmpfs           tmpfs     396M   12K  396M   1% /run/user/42
tmpfs           tmpfs     396M   60K  396M   1% /run/user/0
/dev/vdb1       xfs       10G   33M   10G   1% /gluster02/data1
/dev/vdc1       xfs       10G   33M   10G   1% /gluster02/data2
[root@gluster02 gluster02]#
```

- Gluster fs Package 설치하기 위해 Gluster fs repo 등록
 - # yum install centos-release-gluster



4. 설치 및 실행



4.2 설치 진행(4/5)

```
root@gluster01~# hostname
gluster01
root@gluster01~# yum install centos-release-gluster
Loaded plugins: fastestmirror, langpacks
Loading mirror speeds from cached hostfile
 * base: mirror.kakao.com
 * extras: mirror.kakao.com
 * updates: mirror.kakao.com
Resolving Dependencies
--> Running transaction check
--> Package centos-release-gluster41.noarch 0:1.0-3.el7.centos will be installed
--> Processing Dependency: centos-release >= 7-5.1804.el7.centos.2 for package: centos-relea
se-gluster41-1.0-3.el7.centos.noarch
--> Processing Dependency: centos-release-storage-common for package: centos-release-gluster
41-1.0-3.el7.centos.noarch
--> Running transaction check
--> Package centos-release.x86_64 0:7-5.1804.el7.centos will be updated
--> Package centos-release.x86_64 0:7-5.1804.5.el7.centos will be an update
--> Package centos-release-storage-common.noarch 0:2-2.el7.centos will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package                Arch      Version              Repository           Size
=====
Installing:
centos-release-gluster41      noarch    1.0-3.el7.centos     extras                5.2 k
Installing for dependencies:
centos-release-storage-common noarch    2-2.el7.centos       extras                5.1 k
Updating for dependencies:
centos-release                x86_64    7-5.1804.5.el7.centos updates              25 k
=====

Transaction Summary
-----
Install 1 Package (+1 Dependent package)
Upgrade ( 1 Dependent package)

Total download size: 36 k
Is this ok [y/d/N]: y

Total                               131 kB/s | 36 kB  00:00:00
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction
  Updating : centos-release-7-5.1804.5.el7.centos.x86_64                1/4
  Installing : centos-release-storage-common-2-2.el7.centos.noarch      2/4
  Installing : centos-release-gluster41-1.0-3.el7.centos.noarch        3/4
  Cleanup    : centos-release-7-5.1804.el7.centos.x86_64              4/4
  Verifying  : centos-release-storage-common-2-2.el7.centos.noarch      1/4
  Verifying  : centos-release-gluster41-1.0-3.el7.centos.noarch        2/4
  Verifying  : centos-release-7-5.1804.5.el7.centos.x86_64            3/4
  Verifying  : centos-release-7-5.1804.el7.centos.x86_64              4/4

Installed:
centos-release-gluster41.noarch 0:1.0-3.el7.centos

Dependency Installed:
centos-release-storage-common.noarch 0:2-2.el7.centos

Dependency Updated:
centos-release.x86_64 0:7-5.1804.5.el7.centos

Complete!
root@gluster01~#
```

```
root@gluster02~# hostname
gluster02
root@gluster02~# yum install centos-release-gluster
Loaded plugins: fastestmirror, langpacks
Loading mirror speeds from cached hostfile
 * base: ftp.neowiz.com
 * extras: ftp.neowiz.com
 * updates: ftp.neowiz.com
Resolving Dependencies
--> Running transaction check
--> Package centos-release-gluster41.noarch 0:1.0-3.el7.centos will be installed
--> Processing Dependency: centos-release >= 7-5.1804.el7.centos.2 for package: centos-relea
se-gluster41-1.0-3.el7.centos.noarch
--> Processing Dependency: centos-release-storage-common for package: centos-release-gluste
r41-1.0-3.el7.centos.noarch
--> Running transaction check
--> Package centos-release.x86_64 0:7-5.1804.el7.centos will be updated
--> Package centos-release.x86_64 0:7-5.1804.5.el7.centos will be an update
--> Package centos-release-storage-common.noarch 0:2-2.el7.centos will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package                Arch      Version              Repository           Size
=====
Installing:
centos-release-gluster41      noarch    1.0-3.el7.centos     extras                5.2 k
Installing for dependencies:
centos-release-storage-common noarch    2-2.el7.centos       extras                5.1 k
Updating for dependencies:
centos-release                x86_64    7-5.1804.5.el7.centos updates              25 k
=====

Transaction Summary
-----
Install 1 Package (+1 Dependent package)
Upgrade ( 1 Dependent package)

Total size: 36 k
Total download size: 10 k
Is this ok [y/d/N]: y

Total                               40 kB/s | 10 kB  00:00:00
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction
  Updating : centos-release-7-5.1804.5.el7.centos.x86_64                1/4
  Installing : centos-release-storage-common-2-2.el7.centos.noarch      2/4
  Installing : centos-release-gluster41-1.0-3.el7.centos.noarch        3/4
  Cleanup    : centos-release-7-5.1804.el7.centos.x86_64              4/4
  Verifying  : centos-release-storage-common-2-2.el7.centos.noarch      1/4
  Verifying  : centos-release-gluster41-1.0-3.el7.centos.noarch        2/4
  Verifying  : centos-release-7-5.1804.5.el7.centos.x86_64            3/4
  Verifying  : centos-release-7-5.1804.el7.centos.x86_64              4/4

Installed:
centos-release-gluster41.noarch 0:1.0-3.el7.centos


Dependency Installed:
centos-release-storage-common.noarch 0:2-2.el7.centos

Dependency Updated:
centos-release.x86_64 0:7-5.1804.5.el7.centos

Complete!
root@gluster02~#
```



4. 설치 및 실행



4.2 설치 진행(5/5)

- **Gluster fs Package 설치**

- # yum install gluster*

```
root@gluster01:~# yum install gluster*
Loaded plugins: fastestmirror, langpacks
Loading mirror speeds from cached hostfile
 * base: mirror.kakao.com
 * extras: mirror.kakao.com
 * updates: mirror.kakao.com
Resolving Dependencies
--> Running transaction check
--> Package gluster-block.x86_64 0:0.3-2.e17 will be installed
--> Processing Dependency: tcmu-runner-handler-gifs >= 1.0.4 for package: gluster-block-0.3-2.e17.x86_64
--> Processing Dependency: targetcli >= 2.1.fb43 for package: gluster-block-0.3-2.e17.x86_64
--> Package glusterfs.x86_64 0:3.8.4-53.e17.centos will be updated
--> Package glusterfs.x86_64 0:4.1.5-1.e17 will be an update
--> Package glusterfs-api.x86_64 0:3.8.4-53.e17.centos will be updated
--> Package glusterfs-api.x86_64 0:4.1.5-1.e17 will be an update
libblkid      x86_64      2.23.2-52.e17_5.1    updates      178 k
libmount      x86_64      2.23.2-52.e17_5.1    updates      180 k
libuuid       x86_64      2.23.2-52.e17_5.1    updates      81 k
util-linux    x86_64      2.23.2-52.e17_5.1    updates      2.0 M

Transaction Summary
-----
Install 11 Packages (+20 Dependent packages)
Upgrade 5 Packages (+ 4 Dependent packages)

Total download size: 9.3 M
Is this ok [y/d/N]: y

Dependency Updated:
libblkid.x86_64 0:2.23.2-52.e17_5.1    libmount.x86_64 0:2.23.2-52.e17_5.1
libuuid.x86_64 0:2.23.2-52.e17_5.1      util-linux.x86_64 0:2.23.2-52.e17_5.1

Complete!
root@gluster01:~#
```

```
root@gluster02:~# yum install gluster*
Loaded plugins: fastestmirror, langpacks
Loading mirror speeds from cached hostfile
 * base: ftp.neowiz.com
 * extras: ftp.neowiz.com
 * updates: ftp.neowiz.com
Resolving Dependencies
--> Running transaction check
--> Package gluster-block.x86_64 0:0.3-2.e17 will be installed
--> Processing Dependency: tcmu-runner-handler-gifs >= 1.0.4 for package: gluster-block-0.3-2.e17.x86_64
--> Processing Dependency: targetcli >= 2.1.fb43 for package: gluster-block-0.3-2.e17.x86_64
--> Package glusterfs.x86_64 0:3.8.4-53.e17.centos will be updated
--> Package glusterfs.x86_64 0:4.1.5-1.e17 will be an update
--> Package glusterfs-api.x86_64 0:3.8.4-53.e17.centos will be updated
libmount      x86_64      2.23.2-52.e17_5.1    updates      180 k
libuuid       x86_64      2.23.2-52.e17_5.1    updates      81 k
util-linux    x86_64      2.23.2-52.e17_5.1    updates      2.0 M


Transaction Summary
-----
Install 11 Packages (+20 Dependent packages)
Upgrade 5 Packages (+ 4 Dependent packages)

Total size: 9.3 M
Total download size: 6.9 M
Is this ok [y/d/N]: y

Dependency Updated:
libblkid.x86_64 0:2.23.2-52.e17_5.1    libmount.x86_64 0:2.23.2-52.e17_5.1
libuuid.x86_64 0:2.23.2-52.e17_5.1      util-linux.x86_64 0:2.23.2-52.e17_5.1

Complete!
root@gluster02:~#
```

4. 설치 및 실행



4.3 설치 완료(1/2)

- **Gluster fs Package 설치가 되었는지 확인**
 - Package 설치 확인 : `# rpm -qa | grep gluster`

```
root@gluster01:~# rpm -qa | grep gluster
glusterfs-4.1.5-1.el7.x86_64
glusterfs-api-devel-4.1.5-1.el7.x86_64
glusterfs-rdma-4.1.5-1.el7.x86_64
libvirt-daemon-driver-storage-gluster-3.9.0-14.el7.x86_64
glusterfs-api-4.1.5-1.el7.x86_64
glusterfs-fuse-4.1.5-1.el7.x86_64
glusterfs-server-4.1.5-1.el7.x86_64
glusterfs-resource-agents-4.1.5-1.el7.noarch
glusterfs-libs-4.1.5-1.el7.x86_64
python2-gluster-4.1.5-1.el7.x86_64
glusterfs-cli-4.1.5-1.el7.x86_64
glusterfs-devel-4.1.5-1.el7.x86_64
gluster-block-0.3-2.el7.x86_64
glusterfs-geo-replication-4.1.5-1.el7.x86_64
centos-release-gluster41-1.0-3.el7.centos.noarch
glusterfs-coreutils-0.2.0-1.el7.x86_64
glusterfs-client-xlators-4.1.5-1.el7.x86_64
glusterfs-extra-xlators-4.1.5-1.el7.x86_64
glusterfs-events-4.1.5-1.el7.x86_64
root@gluster01:~#

root@gluster02:~# rpm -qa | grep gluster
glusterfs-4.1.5-1.el7.x86_64
glusterfs-api-devel-4.1.5-1.el7.x86_64
glusterfs-rdma-4.1.5-1.el7.x86_64
libvirt-daemon-driver-storage-gluster-3.9.0-14.el7.x86_64
glusterfs-api-4.1.5-1.el7.x86_64
glusterfs-fuse-4.1.5-1.el7.x86_64
glusterfs-server-4.1.5-1.el7.x86_64
glusterfs-resource-agents-4.1.5-1.el7.noarch
glusterfs-libs-4.1.5-1.el7.x86_64
python2-gluster-4.1.5-1.el7.x86_64
glusterfs-cli-4.1.5-1.el7.x86_64
glusterfs-devel-4.1.5-1.el7.x86_64
gluster-block-0.3-2.el7.x86_64
glusterfs-geo-replication-4.1.5-1.el7.x86_64
centos-release-gluster41-1.0-3.el7.centos.noarch
glusterfs-coreutils-0.2.0-1.el7.x86_64
glusterfs-client-xlators-4.1.5-1.el7.x86_64
glusterfs-extra-xlators-4.1.5-1.el7.x86_64
glusterfs-events-4.1.5-1.el7.x86_64
root@gluster02:~#
```

- **Glusterd daemon 실행 및 Active 상태인지 확인**
 - Glusterd daemon 실행 : `# systemctl start glusterd`
`# systemctl enable glusterd`
 - Glusterfs daemon 실행 확인 : `# systemctl status glusterd`



4. 설치 및 실행



4.3 설치 완료(2/2)

```
root@gluster01:~# systemctl start glusterd
root@gluster01:~# systemctl enable glusterd
Created symlink from /etc/systemd/system/multi-user.target.wants/glusterd.service to /usr/lib/systemd/system/glusterd.service.
root@gluster01:~# systemctl status glusterd
● glusterd.service - GlusterFS, a clustered file-system server
   Loaded: loaded (/usr/lib/systemd/system/glusterd.service; enabled; vendor preset: disabled)
   Active: active (running) since Tue 2018-10-16 18:39:02 KST; 26s ago
     Main PID: 20559 (glusterd)
    CGroup: /system.slice/glusterd.service
            └─20559 /usr/sbin/glusterd -p /var/run/glusterd.pid --log-level INFO

Oct 16 18:39:00 gluster01 systemd[1]: Starting GlusterFS, a clustered file-system server...
Oct 16 18:39:02 gluster01 systemd[1]: Started GlusterFS, a clustered file-system server.
Hint: Some lines were ellipsized, use -l to show in full.
root@gluster01:~#
```

```
root@gluster02:~# systemctl start glusterd
root@gluster02:~# systemctl enable glusterd
Created symlink from /etc/systemd/system/multi-user.target.wants/glusterd.service to /usr/lib/systemd/system/glusterd.service.
root@gluster02:~# systemctl status glusterd
● glusterd.service - GlusterFS, a clustered file-system server
   Loaded: loaded (/usr/lib/systemd/system/glusterd.service; enabled; vendor preset: disabled)
   Active: active (running) since Tue 2018-10-16 18:39:41 KST; 13s ago
     Main PID: 20420 (glusterd)
    CGroup: /system.slice/glusterd.service
            └─20420 /usr/sbin/glusterd -p /var/run/glusterd.pid --log-level INFO

Oct 16 18:39:39 gluster02 systemd[1]: Starting GlusterFS, a clustered file-system server...
Oct 16 18:39:41 gluster02 systemd[1]: Started GlusterFS, a clustered file-system server.
Hint: Some lines were ellipsized, use -l to show in full.
root@gluster02:~#
```



5. 기능소개

세부 목차




1. Volume 신규 구성

1. distributed volume 신규 구성
2. distributed replicated volume 신규 구성



5. 기능소개



5.1.1 distributed volume 신규구성(1/8)

- distributed 구성을 위해 서버 및 brick 확인
 - 2대의 서버
 - gluster01 서버 : 10GB 볼륨(brick) => /gluster01/data1
 - gluster02 서버 : 10GB 볼륨(brick) => /gluster02/data1

```
root@gluster01:~# df -Th
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda3       xfs       21G   3.5G  18G  17% /
devtmpfs        devtmpfs  2.0G   0     2.0G  0% /dev
tmpfs           tmpfs     2.0G   0     2.0G  0% /dev/shm
tmpfs           tmpfs     2.0G   9.2M  2.0G  1% /run
tmpfs           tmpfs     2.0G   0     2.0G  0% /sys/fs/cgroup
/dev/vda1       xfs      1014M  155M  860M  16% /boot
tmpfs           tmpfs     396M   12K  396M  1% /run/user/42
tmpfs           tmpfs     396M   0     396M  0% /run/user/0
/dev/vdb1       xfs       10G    33M  10G  1% /gluster01/data1
/dev/vdc1       xfs       10G    33M  10G  1% /gluster01/data2
root@gluster01:~#
```

```
root@gluster02:~# df -Th
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda3       xfs       21G   3.8G  18G  18% /
devtmpfs        devtmpfs  2.0G   0     2.0G  0% /dev
tmpfs           tmpfs     2.0G   0     2.0G  0% /dev/shm
tmpfs           tmpfs     2.0G   9.4M  2.0G  1% /run
tmpfs           tmpfs     2.0G   0     2.0G  0% /sys/fs/cgroup
/dev/vda1       xfs      1014M  155M  860M  16% /boot
tmpfs           tmpfs     396M   12K  396M  1% /run/user/42
tmpfs           tmpfs     396M   60K  396M  1% /run/user/0
/dev/vdb1       xfs       10G    33M  10G  1% /gluster02/data1
/dev/vdc1       xfs       10G    33M  10G  1% /gluster02/data2
root@gluster02:~#
```

- Storage pool 추가 작업 진행
 - gluster fs 파일시스템을 구성할 서버 등록 및 확인
 - # gluster peer probe gluster01
 - # gluster peer probe gluster02
 - # gluster peer status



5. 기능소개



5.1.1 distributed volume 신규 구성(2/8)

```
root@gluster01:/gluster01/data1
[root@gluster01 data1]# gluster peer probe gluster01
peer probe: success. Probe on localhost not needed.
[root@gluster01 data1]# gluster peer probe gluster02
peer probe: success.
[root@gluster01 data1]# gluster peer status
Number of Peers: 1

Hostname: gluster02
Uid: cf083e93-c9ba-4dcb-9b59-18eb6fb8a38e
State: Peer in Cluster (Connected)
[root@gluster01 data1]#
```

- distributed 구성 작업 진행

- distributed 구성 및 확인

- # gluster volume create distributed-volume gluster01:/gluster01/data1/distvol gluster02:/gluster02/data1/distvol


- # gluster volume info

```
root@gluster01:/gluster01/data1
[root@gluster01 data1]# gluster volume create distributed-volume gluster01:/gluster01/data1/distvol gluster02:/gluster02/data1/distvol
volume create: distributed-volume: success: please start the volume to access data
[root@gluster01 data1]# gluster volume info

Volume Name: distributed-volume
Type: Distribute
Volume ID: 35c865fc-909d-44cb-81ff-9f5bd1ae6c03
Status: Created
Snapshot Count: 0
Number of Bricks: 2
Transport-type: tcp
Bricks:
Brick1: gluster01:/gluster01/data1/distvol
Brick2: gluster02:/gluster02/data1/distvol
Options Reconfigured:
transport.address-family: inet
nfs.disable: on
[root@gluster01 data1]#
```



5. 기능소개




5.1.1 distributed volume 신규구성(3/8)

- gluster volume 사용 가능하도록 전환
gluster volume start distributed-volume

```
[root@gluster01 data1]# gluster volume start distributed-volume  
volume start: distributed-volume: success  
[root@gluster01 data1]# █
```



5. 기능소개




5.1.1 distributed volume 신규구성(4/8)

- glusterfs Volume을 사용하는 서버에서 해당 볼륨을 사용하기 위해 glusterfs repo 등록
yum install centos-release-gluster

```
root@gluster-client:~  
[root@gluster-client ~]# hostname  
gluster-client  
[root@gluster-client ~]# yum install centos-release-gluster  
Loaded plugins: fastestmirror, langpacks  
Loading mirror speeds from cached hostfile  
* base: mirror.navercorp.com  
* extras: mirror.navercorp.com  
* updates: mirror.navercorp.com  
base | 3.6 kB 00:00:00  
extras | 3.4 kB 00:00:00  
updates | 3.4 kB 00:00:00  
updates/7/x86_64/primary_db | 6.0 MB 00:00:00  
Resolving Dependencies  
--> Running transaction check  
--> Package centos-release-gluster41.noarch 0:1.0-3.el7.centos will be installed  
--> Processing Dependency: centos-release >= 7-5.1804.el7.centos.2 for package: centos-rele
```



5. 기능소개



5.1.1 distributed volume 신규 구성(5/8)

```
centos-release-storage-common    noarch    2-2.el7.centos    extras    5.1 k
Updating for dependencies:
centos-release                    x86_64    7-5.1804.5.el7.centos    updates    25 k

Transaction Summary
-----
Install 1 Package (+1 Dependent package)
Upgrade ( 1 Dependent package)

Total size: 36 k
Total download size: 10 k
Is this ok [y/d/N]: y

Installed:
  centos-release-gluster41.noarch 0:1.0-3.el7.centos

Dependency Installed:
  centos-release-storage-common.noarch 0:2-2.el7.centos

Dependency Updated:
  centos-release.x86_64 0:7-5.1804.5.el7.centos

Complete!
[root@gluster-client ~]#
```


- glusterfs Volume을 사용하는 서버에서 glusterfs-client Package 설치

yum install glusterfs-client

```
[root@gluster-client ~]# yum install glusterfs-client
Loaded plugins: fastestmirror, langpacks
Loading mirror speeds from cached hostfile
 * base: mirror.navercorp.com
 * extras: mirror.navercorp.com
 * updates: mirror.navercorp.com
centos-gluster41 | 2.9 kB 00:00:00
centos-gluster41/7/x86_64/primary_db | 77 kB 00:00:00
Resolving Dependencies
```



5. 기능소개



5.1.1 distributed volume 신규 구성(6/8)

```
glusterfs                x86_64                4.1.5-1.el7                centos-gluster41                649 k
glusterfs-api            x86_64                4.1.5-1.el7                centos-gluster41                106 k
glusterfs-cli            x86_64                4.1.5-1.el7                centos-gluster41                203 k
glusterfs-client-xlators x86_64                4.1.5-1.el7                centos-gluster41                970 k
glusterfs-libs           x86_64                4.1.5-1.el7                centos-gluster41                411 k

Transaction Summary
-----
Install 1 Package
Upgrade ( 5 Dependent packages)

Total download size: 2.4 M
Is this ok [y/d/N]: y

Installed:
  glusterfs-fuse.x86_64 0:4.1.5-1.el7

Dependency Updated:
  glusterfs.x86_64 0:4.1.5-1.el7                glusterfs-api.x86_64 0:4.1.5-1.el7
  glusterfs-cli.x86_64 0:4.1.5-1.el7                glusterfs-client-xlators.x86_64 0:4.1.5-1.el7
  glusterfs-libs.x86_64 0:4.1.5-1.el7

Complete!
[root@gluster-client ~]#
```


- glusterfs Volume을 사용하는 서버에서 gluster volume 마운트 진행 및 확인
mkdir /glusterfs
mount -t glusterfs gluster01:distributed-volume /glusterfs
df -h

```
[root@gluster-client ~]# mkdir /glusterfs
[root@gluster-client ~]# mount -t glusterfs gluster01:distributed-volume /glusterfs
[root@gluster-client ~]# df -h
```

| filesystem | Size | Used | Avail | Use% | Mounted on |
|------------------------------|-------|------|-------|------|----------------|
| /dev/vda3 | 21G | 3.9G | 18G | 19% | / |
| devtmpfs | 2.0G | 0 | 2.0G | 0% | /dev |
| tmpfs | 2.0G | 0 | 2.0G | 0% | /dev/shm |
| tmpfs | 2.0G | 9.4M | 2.0G | 1% | /run |
| tmpfs | 2.0G | 0 | 2.0G | 0% | /sys/fs/cgroup |
| /dev/vda1 | 1014M | 155M | 860M | 16% | /boot |
| tmpfs | 396M | 4.0K | 396M | 1% | /run/user/42 |
| tmpfs | 396M | 36K | 396M | 1% | /run/user/0 |
| gluster01:distributed-volume | 20G | 270M | 20G | 2% | /glusterfs |



5. 기능소개



5.1.1 distributed volume 신규구성(7/8)

- glusterfs Volume을 사용하는 서버에서 파일 생성 후 2대의 gluster 서버에서 파일이 분산되어 저장 되는지 확인
 - for 문으로 20개의 파일 생성
 - 서버에서 파일 확인
 - 2대의 gluster에서 파일 확인

```
[root@gluster-client glusterfs]# for ((i=1; i<=20; i++))
> do
> dd if=/dev/zero of=/glusterfs/testfile$i bs=10M count=1
> done
1+0 records in
1+0 records out
10485760 bytes (10 MB) copied, 0.0721611 s, 145 MB/s
1+0 records in
1+0 records out
10485760 bytes (10 MB) copied, 0.0507569 s, 207 MB/s
1+0 records in
1+0 records out
10485760 bytes (10 MB) copied, 0.0490192 s, 214 MB/s
1+0 records in
1+0 records out
10485760 bytes (10 MB) copied, 0.0534076 s, 196 MB/s
1+0 records in
1+0 records out
10485760 bytes (10 MB) copied, 0.0494072 s, 212 MB/s
```



5. 기능소개




5.1.1 distributed volume 신규 구성(8/8)

```
[root@gluster-client glusterfs]# ls
testfile1  testfile11  testfile13  testfile15  testfile17  testfile19  testfile20  testfile4  testfile6  testfile8
testfile10 testfile12  testfile14  testfile16  testfile18  testfile2  testfile3  testfile5  testfile7  testfile9
[root@gluster-client glusterfs]#
```

| | |
|--|--|
| <pre>root@gluster01:/gluster01/data1/distvol [root@gluster01 distvol]# hostname gluster01 [root@gluster01 distvol]# pwd /gluster01/data1/distvol [root@gluster01 distvol]# ls testfile12 testfile16 testfile18 testfile20 testfile7 testfile14 testfile17 testfile19 testfile5 testfile9 [root@gluster01 distvol]#</pre> | <pre>root@gluster02:/gluster02/data1/distvol [root@gluster02 distvol]# hostname gluster02 [root@gluster02 distvol]# pwd /gluster02/data1/distvol [root@gluster02 distvol]# ls testfile1 testfile11 testfile15 testfile3 testfile6 testfile10 testfile13 testfile2 testfile4 testfile8 [root@gluster02 distvol]#</pre> |
|--|--|



5. 기능소개



5.1.2 distributed replicated volume 신규구성(1/8)

- distributed replicated 구성을 위해 서버 및 brick 확인
- 4대의 서버
 - gluster01 서버 : 10GB 볼륨(brick) => /gluster01/data1
 - gluster02 서버 : 10GB 볼륨(brick) => /gluster02/data1
 - gluster03 서버 : 10GB 볼륨(brick) => /gluster03/data1
 - gluster04 서버 : 10GB 볼륨(brick) => /gluster04/data1

```
root@gluster01:gluster01
[root@gluster01 gluster01]# hostname
gluster01
[root@gluster01 gluster01]# df -Th

```

| Filesystem | Type | Size | Used | Avail | Use% | Mounted on |
|------------|----------|-------|------|-------|------|------------------|
| /dev/vda3 | xfs | 21G | 3.5G | 18G | 17% | / |
| devtmpfs | devtmpfs | 2.0G | 0 | 2.0G | 0% | /dev |
| tmpfs | tmpfs | 2.0G | 0 | 2.0G | 0% | /dev/shm |
| tmpfs | tmpfs | 2.0G | 9.2M | 2.0G | 1% | /run |
| tmpfs | tmpfs | 2.0G | 0 | 2.0G | 0% | /sys/fs/cgroup |
| /dev/vda1 | xfs | 1014M | 155M | 860M | 16% | /boot |
| tmpfs | tmpfs | 396M | 12K | 396M | 1% | /run/user/42 |
| tmpfs | tmpfs | 396M | 0 | 396M | 0% | /run/user/0 |
| /dev/vdb1 | xfs | 10G | 33M | 10G | 1% | /gluster01/data1 |
| /dev/vdc1 | xfs | 10G | 33M | 10G | 1% | /gluster01/data2 |

```
[root@gluster01 gluster01]#
```

```
root@gluster02:gluster02
[root@gluster02 gluster02]# hostname
gluster02
[root@gluster02 gluster02]# df -Th


```

| Filesystem | Type | Size | Used | Avail | Use% | Mounted on |
|------------|----------|-------|------|-------|------|------------------|
| /dev/vda3 | xfs | 21G | 3.8G | 18G | 18% | / |
| devtmpfs | devtmpfs | 2.0G | 0 | 2.0G | 0% | /dev |
| tmpfs | tmpfs | 2.0G | 0 | 2.0G | 0% | /dev/shm |
| tmpfs | tmpfs | 2.0G | 9.4M | 2.0G | 1% | /run |
| tmpfs | tmpfs | 2.0G | 0 | 2.0G | 0% | /sys/fs/cgroup |
| /dev/vda1 | xfs | 1014M | 155M | 860M | 16% | /boot |
| tmpfs | tmpfs | 396M | 12K | 396M | 1% | /run/user/42 |
| tmpfs | tmpfs | 396M | 60K | 396M | 1% | /run/user/0 |
| /dev/vdb1 | xfs | 10G | 33M | 10G | 1% | /gluster02/data1 |
| /dev/vdc1 | xfs | 10G | 33M | 10G | 1% | /gluster02/data2 |

```
[root@gluster02 gluster02]#
```



5. 기능소개




5.1.2 distributed replicated volume 신규구성(2/8)

```
root@gluster03:~  
[root@gluster03 ~]# hostname  
gluster03  
[root@gluster03 ~]# df -h  
Filesystem      Size  Used Avail Use% Mounted on  
/dev/vda3       21G   4.0G   18G   19% /  
devtmpfs        2.0G   0     2.0G   0% /dev  
tmpfs           2.0G   0     2.0G   0% /dev/shm  
tmpfs           2.0G   9.2M   2.0G   1% /run  
tmpfs           2.0G   0     2.0G   0% /sys/fs/cgroup  
/dev/vdb1       10G   283M   9.8G   3% /gluster03/data1  
/dev/vdc1       10G    33M   10G    1% /gluster03/data2  
/dev/vda1       1014M 155M   860M  16% /boot  
tmpfs           396M   8.0K   396M   1% /run/user/42  
tmpfs           396M   0     396M   0% /run/user/0  
[root@gluster03 ~]#  
  
root@gluster04:~  
[root@gluster04 ~]# hostname  
gluster04  
[root@gluster04 ~]# df -h  
Filesystem      Size  Used Avail Use% Mounted on  
/dev/vda3       21G   4.0G   18G   19% /  
devtmpfs        2.0G   0     2.0G   0% /dev  
tmpfs           2.0G   0     2.0G   0% /dev/shm  
tmpfs           2.0G   9.2M   2.0G   1% /run  
tmpfs           2.0G   0     2.0G   0% /sys/fs/cgroup  
/dev/vdc1       10G    33M   10G    1% /gluster04/data2  
/dev/vdb1       10G   283M   9.8G   3% /gluster04/data1  
/dev/vda1       1014M 155M   860M  16% /boot  
tmpfs           396M   8.0K   396M   1% /run/user/42  
tmpfs           396M   0     396M   0% /run/user/0  
[root@gluster04 ~]#
```

- Storage pool 추가 작업 진행
 - gluster fs 파일시스템을 구성할 서버를 등록 및 확인
 - # gluster peer probe gluster01
 - # gluster peer probe gluster02
 - # gluster peer probe gluster03
 - # gluster peer probe gluster04
 - # gluster peer status



5. 기능소개



5.1.2 distributed replicated volume 신규구성(3/8)

```
[root@gluster01 ~]# gluster peer probe gluster01
peer probe: success. Probe on localhost not needed
[root@gluster01 ~]# gluster peer probe gluster02
peer probe: success.
[root@gluster01 ~]# gluster peer probe gluster03
peer probe: success.
[root@gluster01 ~]# gluster peer probe gluster04
peer probe: success.
[root@gluster01 ~]# gluster peer status
Number of Peers: 3

Hostname: gluster02
Uuid: 59f8d285-a5da-4b23-90b4-4f641f637471
State: Peer in Cluster (Connected)


Hostname: gluster03
Uuid: 64d8e096-c799-4d9d-b441-5e42391b6f37
State: Peer in Cluster (Connected)

Hostname: gluster04
Uuid: b56b5a9c-4a9e-4396-93e9-dd7a01efd7d9
State: Peer in Cluster (Connected)
[root@gluster01 ~]#
```

- distributed replicated 구성 작업 진행
 - distributed replicated 구성 및 확인
 - # gluster volume create distributed-replicated-volume replica 2 gluster01:/gluster01/data1/drvol gluster02:/gluster02/data1/drvol gluster03:/gluster03/data1/drvol gluster04:/gluster04/data1/drvol
 - # gluster volume info



5. 기능소개




5.1.2 distributed replicated volume 신규구성(4/8)

```
[root@gluster01 ~]# gluster volume create distributed-replicated-volume replica 2 gluster01:/gluster01/data1/drvol gluster02:/gluster02/data1/drvol gluster03:/gluster03/data1/drvol gluster04:/gluster04/data1/drvol
Replica 2 volumes are prone to split-brain. Use Arbiter or Replica 3 to avoid this. See: http://docs.gluster.org/en/latest/Administrator%20Guide/Split%20brain%20and%20ways%20to%20deal%20with%20it/.
Do you still want to continue?
(y/n) y
volume create: distributed-replicated-volume: success: please start the volume to access data
[root@gluster01 ~]# gluster volume info

Volume Name: distributed-replicated-volume
Type: Distributed-Replicate
Volume ID: f487b8af-c6c9-4335-847b-f1bef7a359fe
Status: Created
Snapshot Count: 0
Number of Bricks: 2 x 2 = 4
Transport-type: tcp
Bricks:
Brick1: gluster01:/gluster01/data1/drvol
Brick2: gluster02:/gluster02/data1/drvol
Brick3: gluster03:/gluster03/data1/drvol
Brick4: gluster04:/gluster04/data1/drvol
Options Reconfigured:
transport.address-family: inet
nfs.disable: on
performance.client-io-threads: off
[root@gluster01 ~]# █
```



5. 기능소개




5.1.2 distributed replicated volume 신규구성(5/8)

- gluster volume 사용 가능하도록 전환
gluster volume start distributed-replicated-volume

```
[root@gluster01 ~]#  
[root@gluster01 ~]# gluster volume start distributed-replicated-volume  
volume start: distributed-replicated-volume: success  
[root@gluster01 ~]# █
```



5. 기능소개



5.1.2 distributed replicated volume 신규구성(6/8)

- glusterfs Volume을 사용하는 서버에서 gluster volume 마운트 진행 및 확인

```
# mkdir /glusterfs
```

```
# mount -t glusterfs gluster01:distributed-replicated-volume /glusterfs
```

```
# df -h
```

```
[root@gluster-client ~]#  
[root@gluster-client ~]# mkdir /glusterfs  
[root@gluster-client ~]# mount -t glusterfs gluster01:distributed-replicated-volume /glusterfs  
[root@gluster-client ~]# df -h  
Filesystem                Size      Used Avail Use% Mounted on  
/dev/vda3                  21G       3.9G   18G   19% /  
devtmpfs                   2.0G       0      2.0G    0% /dev  
tmpfs                      2.0G       0      2.0G    0% /dev/shm  
tmpfs                      2.0G     9.1M    2.0G    1% /run  
tmpfs                      2.0G       0      2.0G    0% /sys/fs/cgroup  
/dev/vda1                  1014M    155M    860M   16% /boot  
tmpfs                      396M     12K    396M    1% /run/user/42  
tmpfs                      396M       0    396M    0% /run/user/0  
gluster01:distributed-replicated-volume 20G     270M    20G    2% /glusterfs  
[root@gluster-client ~]#
```



5. 기능소개




5.1.2 distributed replicated volume 신규구성(7/8)

- glusterfs Volume을 사용하는 서버에서 파일 생성 후 4 대의 gluster 서버에 파일이 분산되어 저장 되는지 확인
 - for 문으로 40개의 파일 생성
 - 클라이언트 서버에서 파일 확인
 - 4대의 gluster 서버에서 파일 확인

```
[root@gluster-client ~]# for ((i=1; i<=40; i++))
> do
> dd if=/dev/zero of=/glusterfs/testfile$i bs=10M count=1
> done
1+0 records in
1+0 records out
10485760 bytes (10 MB) copied, 0.0965396 s, 109 MB/s
1+0 records in
1+0 records out
10485760 bytes (10 MB) copied, 0.0744052 s, 141 MB/s
1+0 records in
1+0 records out
10485760 bytes (10 MB) copied, 0.0828328 s, 127 MB/s
1+0 records in
1+0 records out
10485760 bytes (10 MB) copied, 0.119609 s, 87.7 MB/s
1+0 records in
1+0 records out
10485760 bytes (10 MB) copied, 0.0942189 s, 111 MB/s
1+0 records in
1+0 records out
10485760 bytes (10 MB) copied, 0.0979829 s, 107 MB/s
1+0 records in
```



5. 기능소개



5.1.2 distributed replicated volume 신규구성(8/8)

```
[root@gluster-client glusterfs]# ls
testfile1  testfile14  testfile19  testfile23  testfile28  testfile32  testfile37  testfile5
testfile10 testfile15  testfile2  testfile24  testfile29  testfile33  testfile38  testfile6
testfile11 testfile16  testfile20  testfile25  testfile3  testfile34  testfile39  testfile7
testfile12 testfile17  testfile21  testfile26  testfile30  testfile35  testfile4  testfile8
testfile13 testfile18  testfile22  testfile27  testfile31  testfile36  testfile40  testfile9
[root@gluster-client glusterfs]#
```

```
[root@gluster01 drvol]# hostname
gluster01
[root@gluster01 drvol]# pwd
/gluster01/data1/drvol
[root@gluster01 drvol]# ls
testfile12  testfile17  testfile20  testfile30  testfile5
testfile14  testfile18  testfile26  testfile31  testfile7
testfile16  testfile19  testfile27  testfile38  testfile9
[root@gluster01 drvol]#
```

```
[root@gluster02 drvol]# hostname
gluster02
[root@gluster02 drvol]# pwd
/gluster02/data1/drvol
[root@gluster02 drvol]# ls
testfile12  testfile17  testfile20  testfile30  testfile5
testfile14  testfile18  testfile26  testfile31  testfile7
testfile16  testfile19  testfile27  testfile38  testfile9
[root@gluster02 drvol]#
```

```
[root@gluster03 drvol]# hostname
gluster03
[root@gluster03 drvol]# pwd
/gluster03/data1/drvol
[root@gluster03 drvol]# ls
testfile1  testfile2  testfile25  testfile33  testfile39
testfile10 testfile21  testfile28  testfile34  testfile4
testfile11 testfile22  testfile29  testfile35  testfile40
testfile13 testfile23  testfile3  testfile36  testfile6
testfile15 testfile24  testfile32  testfile37  testfile8
[root@gluster03 drvol]#
```

```
[root@gluster04 drvol]# hostname
gluster04
[root@gluster04 drvol]# pwd
/gluster04/data1/drvol
[root@gluster04 drvol]# ls
testfile1  testfile2  testfile25  testfile33  testfile39
testfile10 testfile21  testfile28  testfile34  testfile4
testfile11 testfile22  testfile29  testfile35  testfile40
testfile13 testfile23  testfile3  testfile36  testfile6
testfile15 testfile24  testfile32  testfile37  testfile8
[root@gluster04 drvol]#
```



6. 활용예제

세부 목차



6.1 Gluster fs 구성 방식

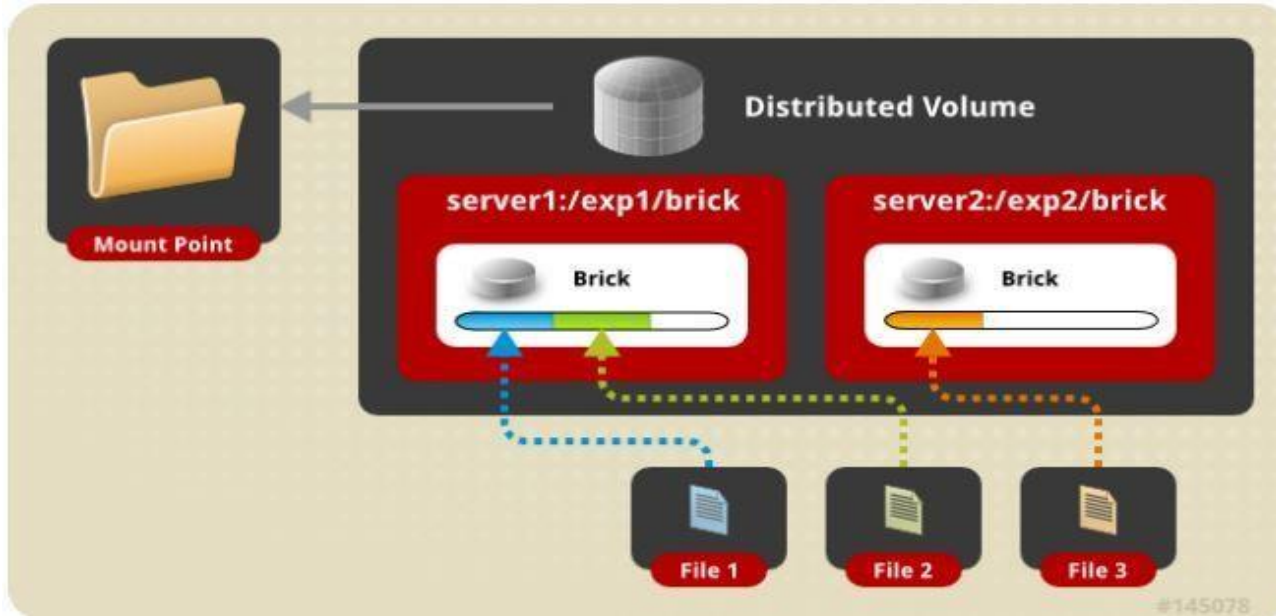


6. 활용예제



6.1 Gluster fs 구성 방식(1/7)


- distributed 구성 방식



- 분산 파일시스템이 필요하고 데이터 안정성이 요구 되지 않을 때 사용
- 여러 개의 brick 볼륨을 거쳐 데이터를 저장하는 방식
- 디스크, node 장애 시 데이터 안전성을 보장 할 수 없음

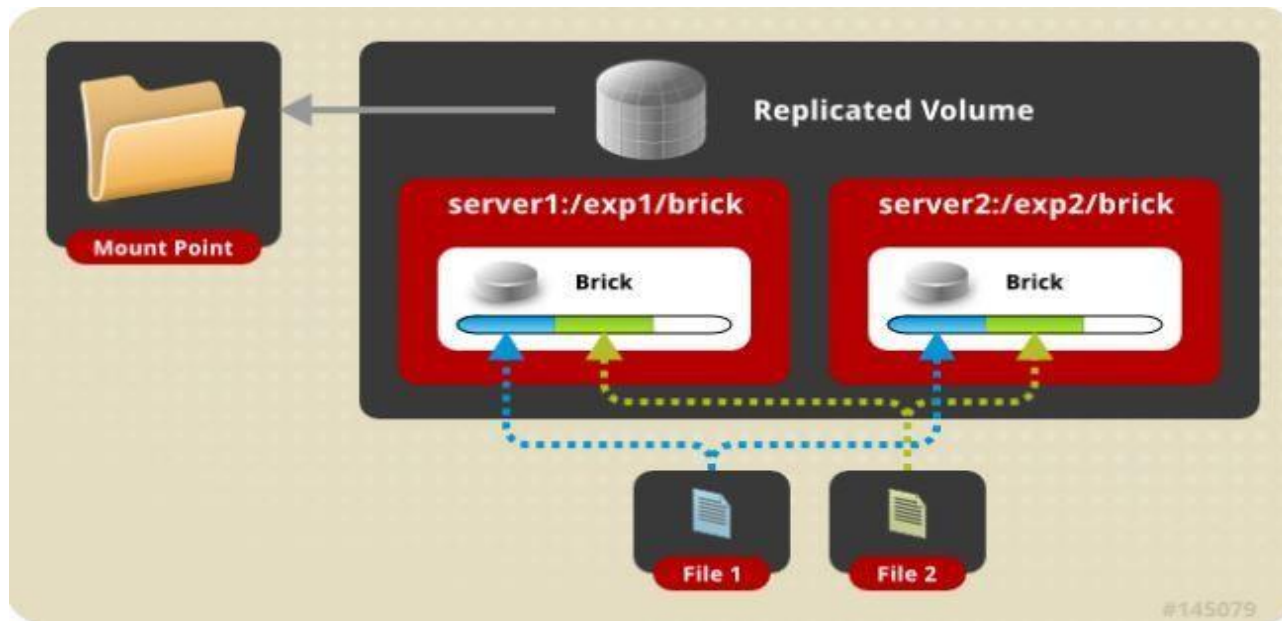


6. 활용예제



6.1 Gluster fs 구성 방식(2/7)


- replicated 구성 방식



- 분산파일 시스템을 사용하지 않고 데이터 안전성이 중요할 때 사용
- 여러 개의 brick 볼륨을 거쳐서 데이터를 저장하는 방식
- 운영하는 데이터의 안전성이 보장되어야 할 때 사용되며, brick 볼륨이 서로 다른 서버에 존재할 때 안전하게 데이터 복제




6. 활용예제



6.1 Gluster fs 구성 방식(3/7)


- striped 구성 방식



- 큰 파일의 동시접속이 많을 때 사용
- 하나의 파일을 여러 개의 brick 볼륨에 분산하여 데이터를 저장 하므로 데이터 안전성을 보장할 수 없음

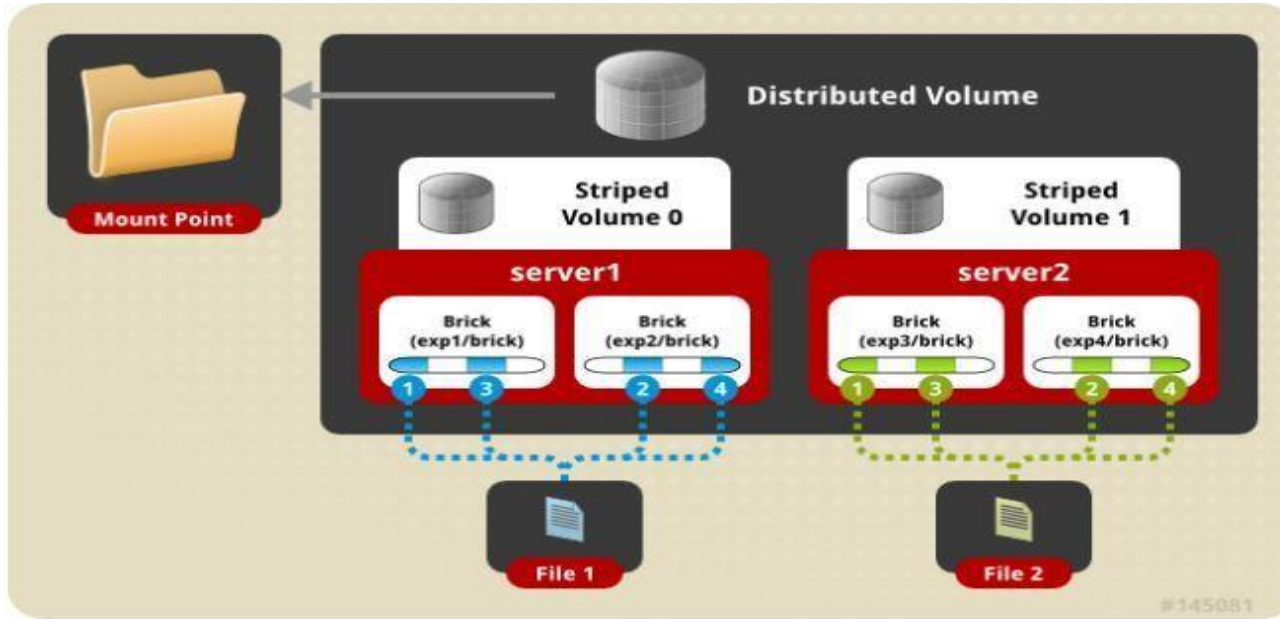


6. 활용예제



6.1 Gluster fs 구성방식(4/7)


- distributed Striped 구성 방식



- 확장성을 요구하며 큰 파일의 동시접속이 많을 때 사용
- 기존의 stripe 방식에서 분산파일 시스템을 사용하기 위해서 구성하는 방식

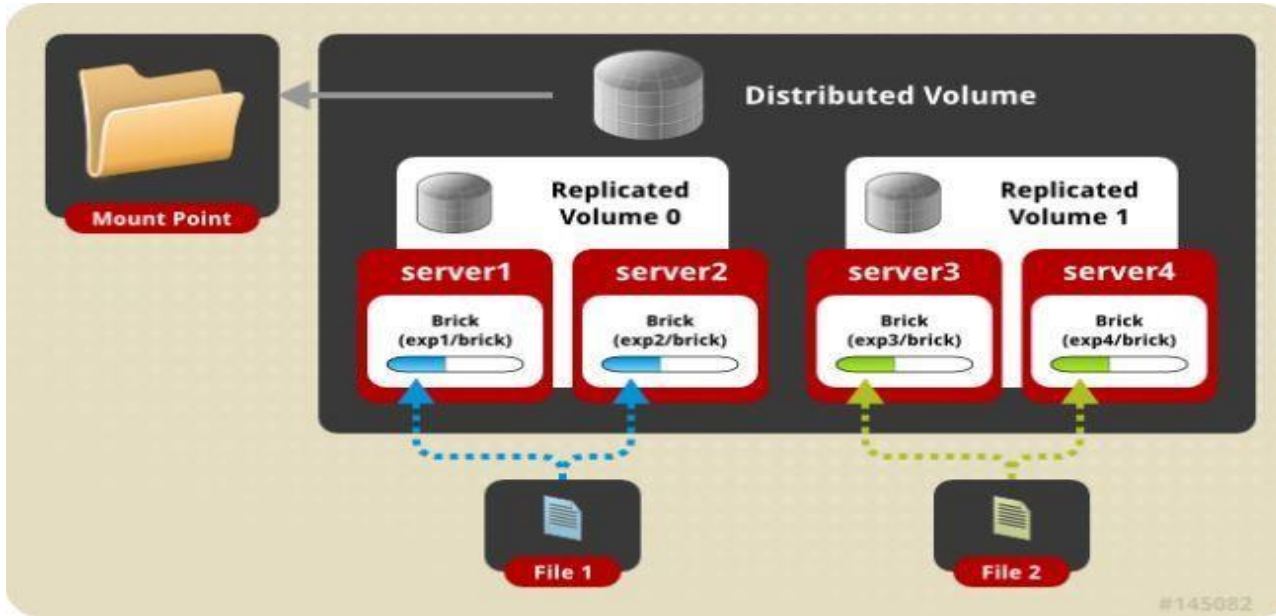


6. 활용예제



6.1 Gluster fs 구성방식(5/7)

- distributed replicated 구성 방식



- 분산파일 시스템을 사용하면서 동시에 데이터 안정성이 중요할 때 사용
- 일반적으로 많이 사용되는 방식으로 Read Performance 좋음

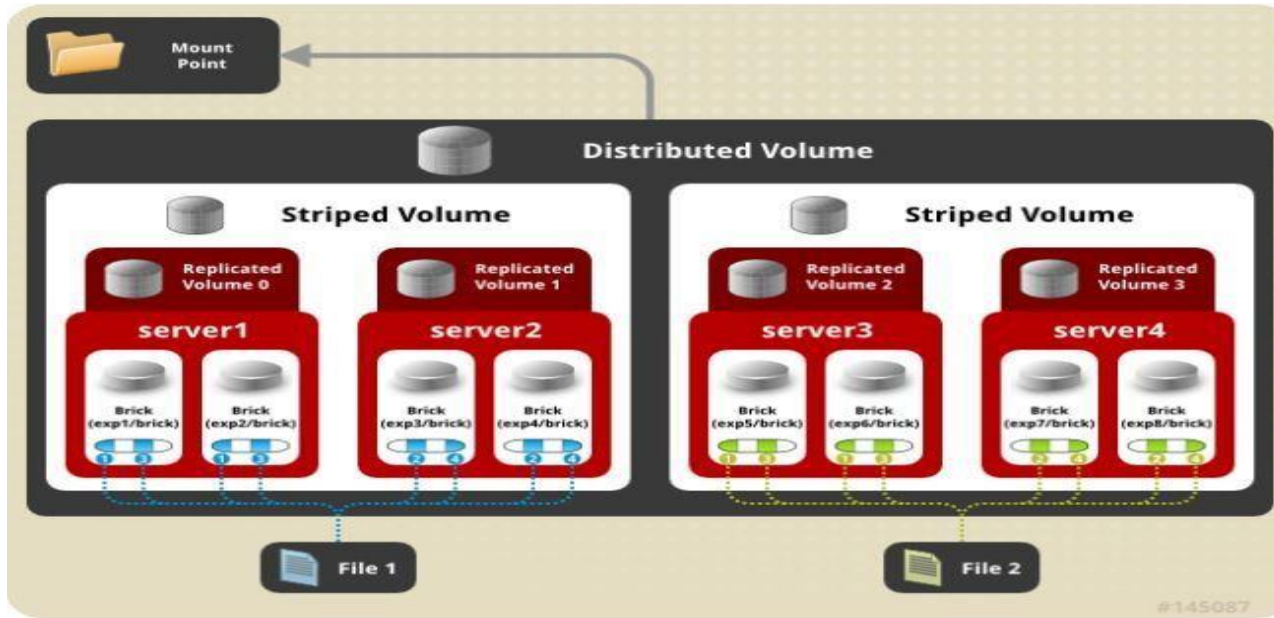


6. 활용예제



6.1 Gluster fs 구성방식(6/7)


- distributed striped replication 구성 방식



- 파일에 대한 동시접속이 많고, 성능이 중요할 때 사용 되는 방식
- map reduce를 사용할 때 구성하는 방식으로 권장
- 복제를 통한 데이터의 안전성을 보장하기 위해서 구성하는 방식




6. 활용예제




6.1 Gluster fs 구성 방식(7/7)

- striped replicated 방식



- 큰 파일에 대한 동시 접속이 많고 성능이 중요할 때 사용
- map reduce를 사용할 때 구성하는 방식으로 권장





Q Gluster fs는 데이터 저장공간에 대한 확장, 축소가 가능 한가요?


A Gluster fs는 처음 구성이 된 이후에도, 데이터 저장공간에 대해 확장 및 축소가 가능 합니다.

Q Gluster fs는 Share Volume으로 사용이 가능 하나요?

A Gluster fs는 NFS(Network File System), SMB Protocol을 지원하므로 Share Volume으로 사용이 가능 하지만, 용도에 맞게 사전 테스트 진행 후 사용하시면 됩니다.



8. 용어정리



| 용어 | 설명 |
|----------|---|
| peer, 노드 | gluster는 분산 파일 시스템 으로 여러 대 노드의 파일시스템을 하나로 묶어서 사용, 이 노드를 'peer' 라고 함 |
| 클러스터 | gluster 에서는 peer 관련 명령어를 통하여 클러스터를 구성 하며, 분산 파일 시스템을 구성 할 peer의 모임을 클러스터 라고 함 |
| brick | brick은 각 노드에서 볼륨으로 사용 할 파일시스템 으로 마운트 포인트를 이용함, brick은 '노드IP:마운트포인트' 의 형태로 표시 (ex - 192.168.1.11:/mnt/test) |
| 볼륨 | 서비스를 제공하는 gluster 볼륨은 여러 개의 brick들로 구성 됨 사용자는 기본적으로 glusterfs 파일시스템을 통해 마운트를 하며, NFS나 SMB, CIFS 등을 통해 마운트 하여 사용 가능 함 |



Open Source Software Installation & Application Guide



이 저작물은 크리에이티브 커먼즈 [저작자표시-비영리-동일조건 변경허락 2.0 대한민국 라이선스]에 따라 이용하실 수 있습니다.