

Configuring an Ignite-UX server under HP Serviceguard



Table of Contents

About this document.....	3
Intended audience	3
Advantage of running Ignite-UX under Serviceguard	3
Setting up Ignite-UX to run under Serviceguard.....	3
Serviceguard packages and scripts	3
Procedure to set up Ignite under Serviceguard	4
1. Create an LVM volume group to manage the shared file systems	4
2. Create logical volumes for the shared file systems	5
3. Create mount points for the shared file systems	6
4. Add the <code>-s</code> option to <code>/usr/sbin/tftpd</code>	7
5. Mount <code>/var/opt/ignite</code> and create required directories on the first cluster node	7
6. Unmount <code>/var/opt/ignite</code> on the first cluster node and deactivate the volume group.....	8
7. Configure group for exclusive access and export	8
8. Import volume group data on all cluster nodes.....	8
9. Copy NFS toolkit cluster scripts in place.....	8
10. Create package config and scripts	9
11. Edit parameters in Ignite scripts and config files on the first node as follows.	9
12. Copy completed package directory to all nodes	11
13. Add the new package to the cluster on the first node.....	11
14. Bring up the package on each node and install Ignite-UX.....	12
15. Update the Ignite server IP address.....	18
16. Create archive directories for cluster nodes	19
17. Add clients and create recovery archives	19
Best practices	19
Upgrading Ignite-UX.....	19
Inactive cluster nodes and <code>swverify</code>	19
Other depots under shared file systems.....	19
Booting from the cluster using <code>/etc/bootptab</code>	20
Edit scripts on all cluster nodes for NFS client mounts	20
When using the Ignite UI or commands, always use the virtual server hostname and IP for the package ...	20
Managing recovery images for cluster nodes.....	20

About this document

This white paper describes the procedure for installing Ignite-UX on an HP Serviceguard cluster running HP-UX 11i v3. It is assumed you have a cluster up and running, there is shared storage available for recovery archives, Serviceguard NFS Toolkit is installed, and Ignite-UX is not installed. This paper covers configuring Ignite-UX to use an LVM shared volume and was written using the September, 2008 DCOE, which contains Serviceguard A.11.18.00 (T1905CA) and Serviceguard NFS Toolkit A.11.31.03 (B5140BA).

Intended audience

This document is intended for system and network administrators responsible for installing, configuring, and managing HP-UX systems. Familiarity with Ignite-UX and HP Serviceguard is assumed. Administrators are expected to have knowledge of operating system concepts, commands, and configuration.

Advantage of running Ignite-UX under Serviceguard

Serviceguard clusters are made up of HP Integrity or HP 9000 servers configured with software and hardware redundancies so your environment continues to run even when there is a failure. Each server in a cluster is called a node.

Serviceguard allows Ignite-UX to run in a clustered environment. One advantage of this is the creation of a highly available recovery server. This setup allows Serviceguard to monitor cluster nodes, networks, and processes, and handle moving the Ignite-UX recovery service to another node in the case of failure.

Setting up Ignite-UX to run under Serviceguard

HP Serviceguard is available as a recommended product in the HP-UX 11i v3 High Availability OE (HA-OE) and the Data Center OE (DC-OE). The Ignite-UX product is included as an optional product on all the HP-UX 11i v3 OEs or can be downloaded via the Ignite-UX product website at: http://h71028.www7.hp.com/enterprise/w1/en/os/hpux11i-system-management-ignite-ux.html?jumpid=ex_r1533_us/en/large/tsg/go_ignite-ux.

For information on how to configure a Serviceguard cluster, see the Serviceguard documentation available at: <http://www.hp.com/go/hpux-serviceguard-docs>

Serviceguard packages and scripts

Serviceguard packages group together applications and the services they depend on. A package and a set of scripts must be created for Ignite. The scripts configure Serviceguard for:

- **Processes to monitor** - for example, NFS is used by Ignite
- **Networks to monitor** - this provides failover if a network interface goes down
- **Storage to manage** – for example, Ignite recovery archives on a RAID array
- **IP addresses to manage** - Ignite "service" has a static IP regardless of the node it is running on

- **File systems to export** - only the active cluster node has exported file systems

Procedure to set up Ignite under Serviceguard

1. Create an LVM volume group to manage the shared file systems

This volume group will contain `/etc/opt/ignite`, `/opt/ignite/boot`, and `/var/opt/ignite`. This must be on a RAID array and must be accessible by each of the other cluster nodes. You must have approximately 1 GB of free disk space available in `/opt/ignite/boot`. For this example, we have a 104 GB disk.

If you use a location other than `/var/opt/ignite` to store archives on the Ignite server, you will need to change these instructions so your custom location is available via NFS as part of this package. If your archives are stored outside the cluster, using a highly available Ignite package in a Serviceguard cluster does not improve the availability of your archives - if the archive server is down, you will not be able to perform any recoveries.

Execute the following on the first cluster node.

```
# diskinfo /dev/rdisk/c13t0d2
SCSI describe of /dev/rdisk/c13t0d2:
    vendor: COMPAQ
    product id: HSV111 (C)COMPAQ
    type: direct access
    size: 104857600 Kbytes
    bytes per sector: 512
#
# mkdir /dev/vg01
# chmod 755 /dev/vg01
# mknod /dev/vg01/group c 64 0x010000
```

For the `mknod` command, the major number is always 64 and the hexadecimal minor number has the form `0xhh0000` where `hh` must be unique to the volume group you are creating. You must use the same minor number for the volume group on all the nodes. You need to determine in advance what minor number is available on all cluster members for this new volume group by reviewing what has been used on all cluster members. Once you determine what minor number is available, use it instead of the `0x010000` value used in the example.

```
# chmod 640 /dev/vg01/group
# pvcreate -f /dev/rdisk/c13t0d2
Physical volume "/dev/rdisk/c13t0d2" has been successfully
created.
# vgcreate -s 16 /dev/vg01 /dev/dsk/c13t0d2
Increased the number of physical extents per physical
volume to 6399.
Volume group "/dev/vg01" has been successfully created.
Volume Group configuration for /dev/vg01 has been saved in
/etc/lvmconf/vg01.conf
```

The name of the volume group is not important; you may use any name supported by Serviceguard instead of `vg01`. If you do use an alternative name, substitute any reference to `vg01` in this documentation with the name you have chosen.

2. Create logical volumes for the shared file systems

You will need logical volumes for `/etc/opt/ignite`, `/var/opt/ignite`, and `/opt/ignite/boot` on the first cluster node.

```
# lvcreate -L 1000 -n lv011 /dev/vg01
Warning: rounding up logical volume size to extent boundary
at size "1008" MB.
Logical volume "/dev/vg01/lv011" has been successfully
created with character device "/dev/vg01/rlv011".
Logical volume "/dev/vg01/lv011" has been successfully
extended.
Volume Group configuration for /dev/vg01 has been saved in
/etc/lvmconf/vg01.conf
# lvcreate -L 100000 -n lv012 /dev/vg01
Logical volume "/dev/vg01/lv012" has been successfully
created with character device "/dev/vg01/rlv012".
Logical volume "/dev/vg01/lv012" has been successfully
extended.
Volume Group configuration for /dev/vg01 has been saved in
/etc/lvmconf/vg01.conf
# lvcreate -L 1000 -n lv013 /dev/vg01
Warning: rounding up logical volume size to extent boundary
at size "1008" MB.
Logical volume "/dev/vg01/lv013" has been successfully
created with character device "/dev/vg01/rlv013".
Logical volume "/dev/vg01/lv013" has been successfully
extended.
Volume Group configuration for /dev/vg01 has been saved in
/etc/lvmconf/vg01.conf
# newfs -F vxfs /dev/vg01/rlv011
    version 6 layout
    1032192 sectors, 1032192 blocks of size 1024, log size
    16384 blocks
    largefiles supported
# newfs -F vxfs /dev/vg01/rlv012
    version 6 layout
    102400000 sectors, 102400000 blocks of size 1024, log
    size 16384 blocks
    largefiles supported
# newfs -F vxfs /dev/vg01/rlv013
    version 6 layout
    1032192 sectors, 1032192 blocks of size 1024, log size
    16384 blocks
    largefiles supported
# lvdisplay /dev/vg01/lv011
--- Logical volumes ---
LV Name                /dev/vg01/lv011
VG Name                /dev/vg01
LV Permission          read/write
LV Status              available/syncd
```

```

Mirror copies                0
Consistency Recovery         MWC
Schedule                     parallel
LV Size (Mbytes)            1008
Current LE                   63
Allocated PE                 63
Stripes                      0
Stripe Size (Kbytes)        0
Bad block                    on
Allocation                   strict
IO Timeout (Seconds)        default
# lvm display /dev/vg01/lvol2
--- Logical volumes ---
LV Name                      /dev/vg01/lvol2
VG Name                      /dev/vg01
LV Permission                read/write
LV Status                    available/syncd
Mirror copies                0
Consistency Recovery         MWC
Schedule                     parallel
LV Size (Mbytes)            100000
Current LE                   6250
Allocated PE                 6250
Stripes                      0
Stripe Size (Kbytes)        0
Bad block                    on
Allocation                   strict
IO Timeout (Seconds)        default
# lvm display /dev/vg01/lvol3
--- Logical volumes ---
LV Name                      /dev/vg01/lvol3
VG Name                      /dev/vg01
LV Permission                read/write
LV Status                    available/syncd
Mirror copies                0
Consistency Recovery         MWC
Schedule                     parallel
LV Size (Mbytes)            1008
Current LE                   63
Allocated PE                 63
Stripes                      0
Stripe Size (Kbytes)        0
Bad block                    on
Allocation                   strict
IO Timeout (Seconds)        default

```

The file systems created above were created with default options. If you require custom options, you may specify them. If your Ignite server will be storing archives in the default location, you must ensure the file system mounted at `/var/opt/ignite` supports largefiles.

3. Create mount points for the shared file systems

You will need `/var/opt/ignite`, `/opt/ignite/boot`, and `/etc/opt/ignite` on all cluster nodes with the correct permissions and ownership.

```

# mkdir -p /etc/opt/ignite
# mkdir -p /var/opt/ignite
# mkdir -p /opt/ignite/boot
# chown bin:bin /var/opt/ignite
# chmod 555 /var/opt/ignite
# chown bin:bin /etc/opt/ignite
# chmod 755 /etc/opt/ignite
# chown bin:bin /opt/ignite/boot
# chmod 555 /opt/ignite/boot

```

4. Add the `-s` option to `/usr/sbin/tftpd`

Add the `-s` option to the `/usr/sbin/tftpd` file on all cluster nodes that can run the Serviceguard package containing Ignite-UX. Each cluster node that can run Ignite-UX must have an entry in `/etc/inetd.conf` that is similar to the following:

```

tftpd      dgram  udp  wait  root  /usr/sbin/tftpd  tftpd -s \
/opt/ignite\
/var/opt/ignite

```

For more information about the `tftpd -s` option, see `tftpd (1m)`.

The directories `/opt/ignite` and `/var/opt/ignite` are also required for Ignite-UX to work correctly. If Ignite-UX is installed on only one cluster node, these directory entries will not be present and must be manually added.

The `/etc/inetd.conf` file must be reconfigured for the changes to take effect. Run the `"inetd -c"` command to reconfigure the `/etc/inetd.conf` file.

If the `-s` option is not present in the `/usr/sbin/tftpd` file, the kernel routing table will determine the IP address used to reply to the client. The booting client might not accept the reply, because this might not be the expected IP address. This might result in intermittent network boot errors. On HP Integrity systems, errors such as PXE-E18 might be displayed, whereas on HP9000 systems, it might be difficult to diagnose the errors. In cases where the `-s` option is not added to the `/usr/sbin/tftpd` file, the possibility of the error occurring depends on package movements and the order in which packages are started in the cluster node on which the Ignite-UX package is running.

5. Mount `/var/opt/ignite` and create required directories on the first cluster node

Create `/var/opt/ignite/data`, `/var/opt/ignite/clients`, and `/var/opt/ignite/recovery/archives` on the first cluster node, then set permissions and ownership. These directories are exported when the service runs, so they must be there prior to installing Ignite.

```

# mount /dev/vg01/lvol2 /var/opt/ignite
# mkdir /var/opt/ignite/data
# mkdir /var/opt/ignite/clients
# mkdir -p /var/opt/ignite/recovery/archives
# chown bin:bin /var/opt/ignite/clients

```

```
# chmod 755 /var/opt/ignite/clients
# chown bin:bin /var/opt/ignite/data
# chmod 755 /var/opt/ignite/data
# chown bin:bin /var/opt/ignite/recovery
# chmod 755 /var/opt/ignite/recovery
# chown bin:bin /var/opt/ignite/recovery/archives
# chmod 755 /var/opt/ignite/recovery/archives
```

6. Unmount `/var/opt/ignite` on the first cluster node and deactivate the volume group

```
# umount /var/opt/ignite
# vgchange -a n vg01
Volume group "vg01" has been successfully changed.
```

7. Configure group for exclusive access and export

Configure `vg01` for exclusive access, then export the volume group to a map file on the first cluster node.

```
# vgchange -c y vg01
Performed Configuration change.
Volume group "vg01" has been successfully changed.
# mkdir -p /etc/cmcluster/vgmap/
# vgexport -p -s -m /etc/cmcluster/vgmap/vg01.map /dev/vg01
vgexport: Preview of vgexport on volume group "/dev/vg01"
succeeded.
```

8. Import volume group data on all cluster nodes

Copy `/etc/cmcluster/vgmap/vg01.map` from the first cluster node to the same location on the remaining cluster nodes, and then run the following commands and `vgimport` on each of the remaining nodes. Make sure to use the same minor number on the remaining nodes that was used initially to set up the first node.

```
# mkdir /dev/vg01
# mknod /dev/vg01/group c 64 0x010000
# vgimport -s -m /etc/cmcluster/vgmap/vg01.map /dev/vg01
vgimport: Volume group "/dev/vg01" has been successfully
created.
Warning: A backup of this volume group may not exist on
this machine.
Please remember to take a backup using the vgcfgbackup
command after activating the volume group.
```

9. Copy NFS toolkit cluster scripts in place

Put all cluster scripts from the NFS toolkit in place on the first node.

```
# mkdir -p /etc/cmcluster/ha-iux
# cp /opt/cmcluster/nfs/* /etc/cmcluster/ha-iux
```


10. Create package config and scripts

Create the package configuration file and the package control script template for Ignite on the first node.

```
# cmmakepkg -p /etc/cmcluster/ha-iux/nfs.conf
The package template has been created.
This file must be edited before it can be used.
# cmmakepkg -s /etc/cmcluster/ha-iux/nfs.cnt1
Package control script is created.
This file must be edited before it can be used.
```

11. Edit parameters in Ignite scripts and config files on the first node as follows.

Substitute your system names for the example system names such as `sgiux3.fc.hp.com` along with your IP addresses.

/etc/cmcluster/ha-iux/hanfs.sh

```
XFS[0]="-o anon=2 /var/opt/ignite/data"
XFS[1]="-o anon=2 /var/opt/ignite/clients"
XFS[2]="-o sec=sys,anon=2,rw=sgiux3.fc.hp.com:
sgiux4.fc.hp.com:sgiux5.fc.hp.com:sgiux6.fc.hp.com
/var/opt/ignite/recovery/archives"
NFS_SERVICE_NAME[0]="nfs.monitor"
NFS_SERVICE_CMD[0]="/etc/cmcluster/ha-iux/nfs.mon"
NFS_SERVICE_RESTART[0]="-r 3"
NFS_FILE_LOCK_MIGRATION=1
```

/etc/cmcluster/ha-iux/nfs.cnt1

```
VG[0]=vg01
LV[0]=/dev/vg01/lvol1; FS[0]=/etc/opt/ignite;
FS_MOUNT_OPT[0]="-o rw"; FS_UMOUNT_OPT[0]="";
FS_FSCK_OPT[0]=""; FS_TYPE[0]="vxfs"
LV[1]=/dev/vg01/lvol2; FS[1]=/var/opt/ignite;
FS_MOUNT_OPT[1]="-o rw"; FS_UMOUNT_OPT[1]="";
FS_FSCK_OPT[1]=""; FS_TYPE[1]="vxfs"
LV[2]=/dev/vg01/lvol3; FS[2]=/opt/ignite/boot;
FS_MOUNT_OPT[2]="-o rw"; FS_UMOUNT_OPT[2]="";
FS_FSCK_OPT[2]=""; FS_TYPE[2]="vxfs"
IP[0]="15.1.53.116"
SUBNET[0]="15.1.48.0"
```

```
function customer_defined_run_cmds
```

```
{
# If /etc/bootptab exists, move aside. Link shared version.
    if [[ -f /etc/bootptab ]];then
        mv /etc/bootptab /etc/bootptab.iux.backup
    fi
    ln -s /etc/opt/ignite/bootptab /etc/bootptab
```

```
# If /etc/dhcptab is not a file or does not exist then we only
# need to create the symlink if it doesn't already exist.
# This can happen in the event of a failure.
```

```

    if [[ ! -f /etc/dhcptab && -f /etc/opt/ignite/dhcptab ]];then
# If it's not a symlink at this point it should be.
        if [[ ! -L /etc/dhcptab ]];then
            rm /etc/dhcptab
            ln -s /etc/opt/ignite/dhcptab /etc/dhcptab
        fi
    fi

# If /etc/dhcptab exists and so does the /etc/opt/ignite/dhcptab
# move the original dhcptab aside and symlink it to the ignite
# one.
    if [[ -f /etc/dhcptab && -f /etc/opt/ignite/dhcptab ]];then
        mv /etc/dhcptab /etc/dhcptab.iux.backup
        ln -s /etc/opt/ignite/dhcptab /etc/dhcptab
    fi

# If recovery_cmds directory is present, assume depot is present
# and register
# NOTE: Additional shared depot registration can be added here
    if [[ -d /var/opt/ignite/depots/recovery_cmds ]];then
        /usr/sbin/swreg -ldepot /var/opt/ignite/depots/recovery_cmds
    fi

    test_return 51
}

function customer_defined_halt_cmds
{

# If /etc/bootptab is a sym link, remove it.
# Restore backup if it exists.
    if [[ -L /etc/bootptab ]];then
        rm /etc/bootptab
    fi

    if [[ -f /etc/bootptab.iux.backup ]];then
        mv /etc/bootptab.iux.backup /etc/bootptab
    fi

# Remove the symlink for /etc/dhcptab if it exists.
    if [[ -L /etc/dhcptab ]];then
        rm /etc/dhcptab
    fi

# Move /etc/dhcptab.iux.backup back to /etc/dhcptab if it exists
    if [[ -f /etc/dhcptab.iux.backup ]];then
        mv /etc/dhcptab.iux.backup /etc/dhcptab
    fi

# If recovery_cmds directory is present, assume depot is present
# and unregister.
# NOTE: Additional shared depots can be handled similarly
    if [[ -d /var/opt/ignite/depots/recovery_cmds ]];then
        /usr/sbin/swreg -ldepot -u /var/opt/ignite/depots/recovery_cmds
    fi
}

```

```
test_return 52
}
```

/etc/cmcluster/ha-iux/nfs.conf

```
PACKAGE_NAME ha-iux
NODE_NAME *
RUN_SCRIPT /etc/cmcluster/ha-iux/nfs.cnt1
HALT_SCRIPT /etc/cmcluster/ha-iux/nfs.cnt1
MONITORED_SUBNET 15.1.48.0
MONITORED_SUBNET_ACCESS FULL
SERVICE_NAME nfs.monitor
SERVICE_FAIL_FAST_ENABLED no
SERVICE_HALT_TIMEOUT 300
```

/etc/cmcluster/ha-iux/nfs.flm

```
NFS_FLM_HOLDING_DIR="/var/opt/ignite/sm"
```

/etc/cmcluster/ha-iux/nfs.mon

```
NFS_FILE_LOCK_MIGRATION=1
```

12. Copy completed package directory to all nodes

Copy package directory `/etc/cmcluster/ha-iux` and its contents to the same location on remaining cluster nodes. At this point, place copies of `/etc/bootptab` and `/etc/dhcptab` under the shared volume `/etc/opt/ignite` on the active cluster node. These are the copies of `bootptab` and `dhcptab` that will be linked to `/etc/bootptab` and `/etc/dhcptab` when the package is started.

13. Add the new package to the cluster on the first node

```
# cmapplyconf -v -P /etc/cmcluster/ha-iux/nfs.conf
Begin package verification...
Checking existing configuration ... Done
Parsing package file: /etc/cmcluster/ha-iux/nfs.conf.
/etc/cmcluster/ha-iux/nfs.conf: A legacy package is being
used.
Attempting to add package ha-iux.
Maximum configured packages parameter is 150.
Configuring 1 package(s).
149 package(s) can be added to this cluster.
200 access policies can be added to this cluster.
Adding the package configuration for package ha-iux.

Modify the package configuration ([y]/n)? y
Completed the cluster update

# cmviewcl

CLUSTER      STATUS
ignite-ux    up
```

NODE	STATUS	STATE
sgiux3	up	running
sgiux4	up	running
sgiux5	up	running
sgiux6	up	running

UNOWNED_PACKAGES

PACKAGE	STATUS	STATE	AUTO_RUN	NODE
ha-iux	down	halted	disabled	unowned

#

14. Bring up the package on each node and install Ignite-UX

Bring up your package on the first node, followed by a `swinstall` of `IGNITE`. Execute `/opt/ignite/lbin/pkg_rec_depot` to create the `recovery_cmds` depot on the first node. Once this is complete, move the package to the remaining nodes one at a time to `swinstall` `IGNITE` on each. Package `ha-iux` must be running on the node to `swinstall` `ignite` correctly.

On the first node, bring up the package and then install Ignite-UX. In this example, the first node is `sgiux3`.

```
# cmrunpkg -n sgiux3 ha-iux
Running package ha-iux on node sgiux3
Successfully started package ha-iux on node sgiux3
cmrunpkg: All specified packages are running

# cmviewcl

CLUSTER      STATUS
ignite-ux    up

NODE         STATUS      STATE
sgiux3      up          running

PACKAGE      STATUS      STATE      AUTO_RUN  NODE
ha-iux       up          running    disabled  sgiux3

NODE         STATUS      STATE
sgiux4      up          running
sgiux5      up          running
sgiux6      up          running

# swinstall -s server:/depot/path IGNITE
===== 02/19/09 14:09:30 MST BEGIN swinstall SESSION
(non-interactive) (jobid=sgiux3-0001)

* Session started for user "root@sgiux3".

* Beginning Selection
```

```

* Target connection succeeded for "sgiux3:/".
* Source connection succeeded for
  "iuxbld:/depot/Ignite-UX/WEB/opt/trunk/curr".
* Source:
  iuxbld:/depot/Ignite-UX/WEB/opt/trunk/curr
* Targets:                sgiux3:/
* Software selections:
  IGNITE,r=C.7.9.90,a=HP-UX_B.11.00_32/64,v=HP
  Ignite-UX.BOOT-COMMON-IA,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.BOOT-COMMON-PA,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.BOOT-KRN-11-11,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.BOOT-KRN-11-23,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.BOOT-KRN-11-31,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.BOOT-SERVICES,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.CD-TOOLS,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.CD-TOOLS-DOC,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.CD-TOOLS-SRC,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.DVD-TOOLS,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.DVD-TOOLS-DOC,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.DVD-TOOLS-SRC,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.FILE-SRV-11-11,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.FILE-SRV-11-23,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.FILE-SRV-11-31,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.IGNITE,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.IGNT-ENG-A-MAN,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.IGNT-JPN-E-MAN,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.IGNT-JPN-S-MAN,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.MGMT-TOOLS,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
  Ignite-UX.RECOVERY,r=C.7.9.90,a=HP-
UX_B.11.11_32/64,v=HP,fr=C.7.9.90,fa=HP-UX_B.11.11_32/64
* Selection succeeded.

* Beginning Analysis and Execution
* Session selections have been saved in the file
  "/.sw/sessions/swinstall.last".
* The analysis phase succeeded for "sgiux3:/".

```

```
* The execution phase succeeded for "sgiux3:/".
* Analysis and Execution succeeded.
```

NOTE: More information may be found in the agent logfile using the command "swjob -a log sgiux3-0001 @ sgiux3:/".

```
===== 02/19/09 14:11:38 MST END swinstall SESSION (non-
interactive)
(jobid=sgiux3-0001)
```

Execute pkg_rec_depot on the first node only:

```
# /opt/ignite/lbin/pkg_rec_depot
* Generating the product specification file.
* Creating depot:
/var/opt/ignite/depots/recovery_cmds

===== 03/12/09 09:40:31 MDT BEGIN swpackage SESSION

* Session started for user "root@sgiux3.fc.hp.com".

* Source:
sgiux3:/var/opt/ignite/tmp/rec_psf.3497
* Target:
sgiux3:/var/opt/ignite/depots/recovery_cmds
* Software selections:
*

* Beginning Selection Phase.
* Reading the Product Specification File (PSF)
"/var/opt/ignite/tmp/rec_psf.3497".
* Reading the product "Ignite-UX" at line 5.
* Reading the fileset "BOOT-COMMON-IA" at line 20.
* Reading the fileset "BOOT-COMMON-PA" at line 46.
* Reading the fileset "BOOT-KRN-11-11" at line 71.
* Reading the fileset "BOOT-KRN-11-23" at line 106.
* Reading the fileset "BOOT-KRN-11-31" at line 144.
* Reading the fileset "BOOT-SERVICES" at line 182.
* Reading the fileset "CD-TOOLS" at line 214.
* Reading the fileset "CD-TOOLS-DOC" at line 248.
* Reading the fileset "CD-TOOLS-SRC" at line 261.
* Reading the fileset "DVD-TOOLS" at line 275.
* Reading the fileset "DVD-TOOLS-DOC" at line 297.
* Reading the fileset "DVD-TOOLS-SRC" at line 312.
* Reading the fileset "FILE-SRV-11-11" at line 325.
* Reading the fileset "FILE-SRV-11-23" at line 351.
* Reading the fileset "FILE-SRV-11-31" at line 380.
* Reading the fileset "IGNITE" at line 405.
* Reading the fileset "IGNT-ENG-A-MAN" at line 559.
* Reading the fileset "IGNT-JPN-E-MAN" at line 612.
* Reading the fileset "IGNT-JPN-S-MAN" at line 663.
* Reading the fileset "MGMT-TOOLS" at line 714.
* Reading the fileset "RECOVERY" at line 866.
```

```

NOTE:      Creating new target depot
           "/var/opt/ignite/depots/recovery_cmds".
* Selection Phase succeeded.

* Beginning Analysis Phase.
* Analysis Phase succeeded.

* Beginning Package Phase.
NOTE:      You have specified the "package_in_place" option.
The software you have selected will be packaged with
references to the actual source files.  No source files
will be copied into the target depot
"/var/opt/ignite/depots/recovery_cmds".
* Packaging the product "Ignite-UX".
* Packaging the fileset "Ignite-UX.BOOT-COMMON-IA".
* Packaging the fileset "Ignite-UX.BOOT-COMMON-PA".
* Packaging the fileset "Ignite-UX.BOOT-KRN-11-11".
* Packaging the fileset "Ignite-UX.BOOT-KRN-11-23".
* Packaging the fileset "Ignite-UX.BOOT-KRN-11-31".
* Packaging the fileset "Ignite-UX.BOOT-SERVICES".
* Packaging the fileset "Ignite-UX.CD-TOOLS".
* Packaging the fileset "Ignite-UX.CD-TOOLS-DOC".
* Packaging the fileset "Ignite-UX.CD-TOOLS-SRC".
* Packaging the fileset "Ignite-UX.DVD-TOOLS".
* Packaging the fileset "Ignite-UX.DVD-TOOLS-DOC".
* Packaging the fileset "Ignite-UX.DVD-TOOLS-SRC".
* Packaging the fileset "Ignite-UX.FILE-SRV-11-11".
* Packaging the fileset "Ignite-UX.FILE-SRV-11-23".
* Packaging the fileset "Ignite-UX.FILE-SRV-11-31".
* Packaging the fileset "Ignite-UX.IGNITE".
* Packaging the fileset "Ignite-UX.IGNT-ENG-A-MAN".
* Packaging the fileset "Ignite-UX.IGNT-JPN-E-MAN".
* Packaging the fileset "Ignite-UX.IGNT-JPN-S-MAN".
* Packaging the fileset "Ignite-UX.MGMT-TOOLS".
* Packaging the fileset "Ignite-UX.RECOVERY".
* Package Phase succeeded.

NOTE:      You must register the new depot
           "/var/opt/ignite/depots/recovery_cmds" to make it
generally available as a source for swinstall and swcopy
tasks.  To register it, execute the command

           swreg -l depot
/var/opt/ignite/depots/recovery_cmds

===== 03/12/09 09:40:39 MDT  END swpackage SESSION

* Registering depot:
/var/opt/ignite/depots/recovery_cmds

===== 03/12/09 09:40:39 MDT  BEGIN swreg SESSION (non-
interactive)

```

```

* Session started for user "root@sgiux3".

* Beginning Selection
* Targets:                sgiux3
* Objects:
/var/opt/ignite/depots/recovery_cmds
* Selection succeeded.

===== 03/12/09 09:40:39 MDT  END swreg SESSION (non-
interactive)

* Creating IUX Recovery bundles.

===== 03/12/09 09:40:40 MDT  BEGIN swpackage SESSION

* Session started for user "root@sgiux3.fc.hp.com".

* Source:                sgiux3:/var/opt/ignite/tmp/psf.3579
* Target:
sgiux3:/var/opt/ignite/depots/recovery_cmds
* Software selections:
    *

* Beginning Selection Phase.
* Reading the Product Specification File (PSF)
  "/var/opt/ignite/tmp/psf.3579".
* Reading the bundle "Ignite-UX-11-11" at line 11.

* Selection Phase succeeded.

* Beginning Analysis Phase.
* Analysis Phase succeeded.

* Beginning Package Phase.

* Packaging the bundle
  "Ignite-UX-11-11,r=C.7.8.201,a=HP-
UX_B.11.11_32/64,v=HP".
* Package Phase succeeded.

===== 03/12/09 09:40:41 MDT  END swpackage SESSION

===== 03/12/09 09:40:41 MDT  BEGIN swpackage SESSION

* Session started for user "root@sgiux3.fc.hp.com".

* Source:                sgiux3:/var/opt/ignite/tmp/psf.3605
* Target:
sgiux3:/var/opt/ignite/depots/recovery_cmds
* Software selections:

```



```

*
* Beginning Selection Phase.
* Reading the Product Specification File (PSF)
  "/var/opt/ignite/tmp/psf.3605".
* Reading the bundle "Ignite-UX-11-23" at line 11.
* Selection Phase succeeded.
* Beginning Analysis Phase.
* Analysis Phase succeeded.
* Beginning Package Phase.
* Packaging the bundle
  "Ignite-UX-11-23,r=C.7.8.201,a=HP-
UX_B.11.11_32/64,v=HP".
* Package Phase succeeded.
===== 03/12/09 09:40:41 MDT  END swpackage SESSION
===== 03/12/09 09:40:41 MDT  BEGIN swpackage SESSION
* Session started for user "root@sgiux3.fc.hp.com".
* Source:          sgiux3:/var/opt/ignite/tmp/psf.3631
* Target:
sgiux3:/var/opt/ignite/depots/recovery_cmds
* Software selections:
  *
* Beginning Selection Phase.
* Reading the Product Specification File (PSF)
  "/var/opt/ignite/tmp/psf.3631".
* Reading the bundle "Ignite-UX-11-31" at line 11.
* Selection Phase succeeded.
* Beginning Analysis Phase.
* Analysis Phase succeeded.
* Beginning Package Phase.
* Packaging the bundle
  "Ignite-UX-11-31,r=C.7.8.201,a=HP-
UX_B.11.11_32/64,v=HP".
* Package Phase succeeded.
===== 03/12/09 09:40:41 MDT  END swpackage SESSION

```

```

===== 03/12/09 09:40:42 MDT BEGIN swpackage SESSION

* Session started for user "root@sgiux3.fc.hp.com".

* Source:          sgiux3:/var/opt/ignite/tmp/psf.3664
* Target:
sgiux3:/var/opt/ignite/depots/recovery_cmds
* Software selections:
      *

* Beginning Selection Phase.
* Reading the Product Specification File (PSF)
  "/var/opt/ignite/tmp/psf.3664".
* Reading the bundle "IUX-Recovery" at line 11.

* Selection Phase succeeded.

* Beginning Analysis Phase.
* Analysis Phase succeeded.

* Beginning Package Phase.

* Packaging the bundle
  "IUX-Recovery,r=C.7.8.201,a=HP-
UX_B.11.11_32/64,v=HP".
* Package Phase succeeded.

===== 03/12/09 09:40:42 MDT END swpackage SESSION

```

Now move the package to the remaining nodes one at a time to `swinstall IGNITE` on each.

15. Update the Ignite server IP address

Update the Ignite server IP address and execute optional `setup_server` customizations on the cluster node where the `ha-iux` service is running. Note that the Ignite server IP address defaults to the physical IP address of the last cluster node that Ignite was installed on. This needs to be updated to the virtual IP address assigned to the Serviceguard package for Ignite.

```
# /opt/ignite/bin/instl_admin -t [ha-iux virtual hostname or IP]
```

The `setup_server` command may be executed here. Note that the `setup_server` option `(-n)` for NFS mounts does not apply, as these are managed by Serviceguard using the entries from the control scripts.

16. Create archive directories for cluster nodes

Log into the virtual service for the package and create archive directories for all cluster nodes.

```
# cd /var/opt/ignite/recovery/archives
# mkdir -p sgiux3 sgiux4 sgiux5 sgiux6
# chown bin:bin *
# cd /
```

17. Add clients and create recovery archives

Add all cluster nodes as clients and then create recovery archives for each inactive node (not running the `ha-iux` package). From each node, execute the following:

```
# /opt/ignite/bin/add_new_client \
-s [ha-iux virtual hostname or IP]
```

From each inactive node, create a recovery archive. These can run in parallel using the virtual server IP. When they all complete, the package can be moved to another node to allow archive creation for the last node.

```
# /opt/ignite/bin/make_net_recovery -x inc_entire=vg00 \
-s [ha-iux virtual hostname or IP]
```

Best practices

Upgrading Ignite-UX

To upgrade the version of Ignite-UX on the cluster, you will need to `swinstall` the newer version on each node one at a time. The Serviceguard package must be running on the node that is updating. This `swinstall` will execute `pkg_rec_depot` when necessary. After the last node is updated, the customizations from step 15 will need to be re-applied.

Inactive cluster nodes and `swverify`

On nodes that are not running the Ignite Serviceguard package, `swverify` of IGNITE will fail. This is expected due to `/var/opt/ignite`, `/etc/opt/ignite`, and `/opt/ignite/boot` being set up for exclusive access.

Other depots under shared file systems

If there are other SD depots under shared file systems such as `/var/opt/ignite`, they will need to be registered on each cluster node individually while the Ignite package is running on that node, or registered by Serviceguard when the Ignite package starts. To manage from Serviceguard, The `swreg` commands can be added to the `/etc/cmcluster/ha-iux/nfs.cntl` `customer_defined_run_cmds` and `customer_defined_halt_cmds` similar to the handling of the `/var/opt/ignite/depots/recovery_cmds` depot. When you `swinstall` from these depots, be sure to use the virtual hostname or IP for the Ignite Serviceguard package.

Booting from the cluster using `/etc/bootptab`

All `bootptab` entries must be maintained on the active cluster node, as this copy is linked to the shared version `/etc/opt/ignite/bootptab`. The local copies of `/etc/bootptab` on the inactive cluster nodes must not contain entries for clients that will be booting from the Ignite server.

Edit scripts on all cluster nodes for NFS client mounts

The `/etc/dfs/dfstab` file is not used in this configuration; the Serviceguard control scripts will export all file systems. Edit all cluster nodes' control scripts to include NFS exports for recovery clients. In our case, `/etc/cmcluster/ha-iux/hanfs.sh` (from the NFS toolkit template) has the following:

```
# NFS: Specifies all exported NFS directories with share options. The
# directories must be defined in the mounted file system FS list
specified
# in the control script. These variables are used by the command
"share"
# to share the file systems and the command "unshare" to unshare the
filesystems.
# Example: XFS[0]="-o ro /pkg1"
# Example: XFS[1]="/pkg2"
#XFS[0]=" "
XFS[0]="-o anon=2 /var/opt/ignite/data"
XFS[1]="-o anon=2 /var/opt/ignite/clients"
XFS[2]="-o
sec=sys,anon=2,rw=sgiux3.fc.hp.com:sgiux4.fc.hp.com:sgiux5.fc.hp.com:
sgiux6.fc.hp.com /var/opt/ignite/recovery/archives"
```

When using the Ignite UI or commands, always use the virtual server hostname and IP for the package

If you don't, you will have problems when running on other nodes, since Ignite stores the server IP in its config files.

Managing recovery images for cluster nodes

Recovery images for the individual cluster nodes can be created and restored as well, providing at least one cluster node remains running throughout the process. Always do this for nodes that are inactive. For example, if a package is running on `node1`, you can run `make_net_recovery` on `node2` and also re-ignite `node2` from the cluster itself. Again, always use the virtual server hostname and IP address.

For more information

The following relevant documents are available online at the HP Technical Documentation website at [HP Business Support Center](#):

Managing Serviceguard

Ignite-UX Administration Guide

HP-UX 11i v3 Release Notes

HP-UX System Administrator's Guide

Some or all of these documents are available on the Instant Information media and in printed form.

Product information regarding Ignite-UX is available at:

<http://www.hp.com/go/ignite-ux-docs>

© Copyright 2009, 2011 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

UNIX is a registered trademark of The Open Group.

Intel Itanium® Logo, Intel, Intel Inside and Itanium are trademarks or registered trademarks of Intel Corporation in the US and other countries and are used under license.

Intel® Itanium® Processor Family is a trademark of Intel Corporation in the US and other countries and is used under license

MPN 5992-5840, 10/2011