

Tru64 UNIX Best Practice

Configuring a Dual-Alias Cluster on Public and Private Subnets

This Best Practice describes how to configure a dual-alias cluster on public and private subnets for TruCluster Server Version 5.1A and later.

Contents

Configuring Dual-Alias Cluster on Public and Private Subnets

Is This Best Practice Right for You?	1
What Problem Does This Best Practice Address?	2
Before You Begin	3
Applying the Best Practice	5
Configuring a Dual-Alias Cluster When Creating a Cluster	5
Configuring a Dual-Alias Cluster After a Cluster Is Created	7
Monitoring the Private Subnet with CAA	8
Verifying Success	8
Troubleshooting	9
Comments and Questions	9
Legal Notice	9

Configuring Dual-Alias Cluster on Public and Private Subnets

This Best Practice describes how to isolate one or more TruCluster™ Server cluster members to a private subnet while also retaining some connections with the public subnet for the Tru64™ UNIX operating system.

See the Tru64 UNIX Best Practices Web page for more information about Best Practices documentation:

http://www.tru64unix.compaq.com/docs/best_practices/

Is This Best Practice Right for You?

Not all Best Practices apply to all configurations, so you must be sure that this Best Practice is appropriate for your system and circumstances. To use this Best Practice, you must meet the requirements described in the following table:

Requirement	Description
Operating System	Tru64 UNIX Version 5.1A or higher and TruCluster Server Version 5.1A or higher.
Impact on Availability	This Best Practice does not affect system availability.
Skill Level Requirements	This Best Practice is intended for experienced system administrators. In particular, you must be familiar with the process of installing and configuring a TruCluster Server as described in the first five chapters of the <i>Cluster Installation</i> manual. For the purpose of keeping the information described here as concise as possible, this Best Practice summarizes information and procedures described more completely in those chapters.

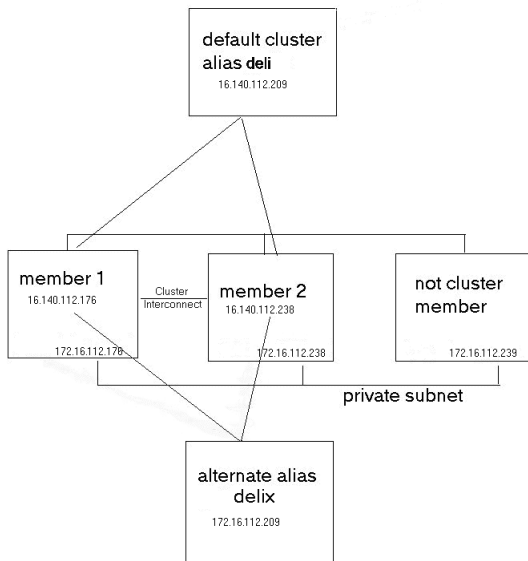
What Problem Does This Best Practice Address?

In certain situations you might want to isolate one or more cluster members on a private subnet while at the same time retaining connections with the public subnet. For example, you might want to eliminate network traffic between the public subnet and the designated cluster members while at the same time using a display or network file system (NFS) server system on a public subnet.

By assigning private subnet addresses to network interfaces and creating an alternate cluster alias by which the cluster members can be addressed, a cluster can eliminate network traffic between the public subnet and the designated cluster members, restrict access, and so forth.

Consider the example shown in *Dual-Alias Cluster on Public and Private Subnets*.

Dual-Alias Cluster on Public and Private Subnets



Member 1 and Member 2 respond to the `deli` alias and to the `delix` alias. The `delix` cluster alias listens for client requests directed to the `delix` alias and distributes them in a round-robin fashion between services running on the two members.

In this example, Member 1 and Member 2 run an HTTP server configured under the `delix` cluster alias on the private subnet. This HTTP

server is not aware of the public subnet. The HTTP client runs on the non-cluster-member system and addresses the HTTP server only via the `delix` alias. Therefore, all HTTP traffic to and from this HTTP server uses only the private subnet.

This configuration also limits access to the HTTP server and the data that it serves. Only those systems on the private subnet are aware of and can access the `delix` alias and, therefore, the HTTP server and its data.

As described in the Tru64 UNIX *Network Administration: Connections* manual, you can configure multiple active network interfaces in one system. You can do so whether the network interfaces operate on the same subnet (for example, a `tu0` interface at 16.1.1.1 and a `tu1` interface at 16.1.1.2) or whether the network interfaces operate in different subnets.

You can think of a cluster alias as a distributed virtual clusterwide network interface. The cluster alias subsystem allows TCP and UDP applications to address the cluster as though it were a single system. When the cluster is created, a default cluster alias is defined that addresses all cluster members. All members can receive packets that are addressed to this alias.

A site can then define additional aliases that address some or all cluster members. Each system in a cluster explicitly joins the aliases to which it wants to belong. After a system joins an alias, it is a member of that alias. Using the analogy that a cluster alias is similar to an address on a virtual network interface, joining an alias is similar to issuing an `ifconfig up` command for that alias interface. The member can now receive packets addressed to the alias.

Several standard Internet services, such as `telnet` and `login`, use the IP address of the default cluster alias as the source address for outgoing packets. Therefore, do not use the default cluster alias as the alias on the private subnet.

Before You Begin

This Best Practice describes how to configure a dual-alias cluster on public and private subnets. The Best Practice assumes that you have read and are familiar with the procedures described in the *Cluster Installation* manual for installing and configuring TruCluster Server. Tables 2-1, 3-1, 4-1, and 5-1 of that manual summarize the preparation tasks and reference sources of information that provide more detail.

You also need the following information:

- A cluster alias name for the private subnet.
- For each network interface to be used on the private subnet, the fully qualified host name and IP address you will use when configuring that interface.

To make it easier to identify the actual physical systems, we recommend that the host name on the private subnet identify the host name used by the interface on the public network. For example, if the host name of the system on the public network is `polishham.zk3.dec.com`, you might use `polishhamx.zk3.dec.com` to identify the interface on the private subnet.

The *Network Administration: Connections* manual provides guidelines on allocating IP addresses. See RFC 1918, in which the Internet Assigned Numbers Authority (IANA) reserves the following blocks of IP address space for use by private internets:

```
10.0.0.0    - 10.255.255.255
172.16.0.0 - 172.31.255.255
192.168.0.0 - 192.168.255.255
```

The sample configuration discussed in this Best Practice consists of a two-node cluster known on a public subnet as `deli` and on a private subnet as `delix`.

The `deli` cluster consists of two members: `polishham` and `provolone`. The `delix` cluster also consists of two members: `polishhamx` and `provolonex`. Physically, `polishhamx` is the same system as `polishham` and `provolonex` is the same system as `provolone`.

All systems belong to the same BIND domain, `zk3.dec.com`. The public subnet IP addresses are selected to be `16.140.112.xxx`. The private subnet IP addresses are selected to be `172.16.112.xxx`, where `xxx` for each physical system is the same on the public and private subnet. The `172.16.112.xxx` addresses were selected to avoid a conflict with the `10.` addresses typically used for the cluster interconnect.

As described in `cluamgr(8)`, by default the `aliasd` daemon does not advertise alias IP addresses that are configured in any of the reserved private address spaces defined in RFC 1918 (such as `10.0.0.0`, `172.16.0.0`, and `192.168.0.0`). That is, unless you specify the `resvok` option, `aliasd` does not advertise network or host routes to alias addresses in the private subnet described in this Best Practice.

If you do specify the `resvok` option, `aliasd` advertises alias IP addresses that reside in reserved subnets, and also advertises alias IP addresses on network interfaces configured within these reserved subnets.

This Best Practice assumes that you do not specify the `resvok` option because you do not want the systems on the private subnet to be available from the public subnet.

Applying the Best Practice

You can configure a dual-alias cluster when you create a cluster, as described in *Configuring a Dual-Alias Cluster When Creating a Cluster*, or anytime after the cluster is created, as described in *Configuring a Dual-Alias Cluster After a Cluster Is Created*.

Before you configure a dual-alias cluster on public and private subnets, be sure to follow the recommendations in *Before You Begin*.

Configuring a Dual-Alias Cluster When Creating a Cluster

Perform the following steps:

1. Install the Tru64 UNIX operating system software on the system that will serve as the first member of the cluster, as described in Chapter 3 of the *Cluster Installation* manual.
2. In addition to the network interface that you want to use on the public network, configure a network interface to use on the private subnet.

Assign a host name and use an IP address reserved for private subnets. The example uses a host name that reflects the host name of the interface on the public network. The host portion of the IP address is the same as that used on the public network to make it easy to identify.

```
interface type:          tu0
host name for this interface:  provolonex.zk3.dec.com
IP address for this interface: 172.16.112.238
```

3. Configure BIND, NFS, NTP, and set up services as described in the *Cluster Installation* manual.
4. If you are installing a patch kit as part of the base operating system installation, load the TruCluster Server subsets before installing the patch kit. If the TruCluster Server kit is not loaded before the patch operation, patches for TruCluster Server software will not be loaded.
5. Create a single-member cluster and boot the system as described in Chapter 4 of the *Cluster Installation* manual.

6. Add and boot another cluster member as described in Chapter 5 of the *Cluster Installation* manual.
7. Configure the system network interface on the added member that you want to use on the private subnet.

During its first boot, the added member automatically runs a script so you can configure additional network interfaces. (The `clu_add_member` command configures only the cluster interconnect interface.) You need to configure at least two interfaces: the public network interface and the private subnet interface.

Assign a host name and use an IP address reserved for private subnets. The example uses a host name that reflects the host name of the interface on the public network. The host portion of the IP address is the same as that used on the public network to make it easy to identify.

```
interface type:          ee2
host name for this interface:  polishhamx.zk3.dec.com
IP address for this interface: 172.16.112.176
```

8. Reboot the added member so that it is using its custom kernel, and perform additional add-member tasks as described in Chapter 5 of the *Cluster Installation* manual.
9. Repeat steps 6 through 8 to add any other members.
10. Identify the members that will constitute the cluster on the private subnet and configure their network interfaces.
11. Edit the `/etc/hosts` file and add the cluster alias for the private subnet. Specify the IP address of the cluster alias, the fully qualified host name including the domain name, and the unqualified host name for this cluster alias:

```
172.16.112.209    delix.zk3.dec.com    delix
```

When all members have been added to the cluster, identify those members that will constitute the cluster on the private subnet. Then, on each such member, perform the following steps:

1. Join the cluster alias you want to use for the cluster on the private subnet. To do this, execute the following command on each member that you want to join the alias:
- ```
cluamgr -a alias=delix.zk3.dec.com,join
```
2. On each member that you want to join the alias, edit the `/etc/clu_alias.config` file and add an entry for the cluster alias on the private subnet:

```
/usr/sbin/cluamgr -a selw=3,selp=1,join,alias=delix.zk3.dec.com
```

The `/etc/clu_alias.config` file is a context-dependent symbolic link (CDSL) pointing to member-specific cluster alias configuration files. (The `clu_create` and `clu_add_member` commands add the default cluster entry to the file.) See `cluamgr(8)` for a description of the `selw` and `selp` options.

## Configuring a Dual-Alias Cluster After a Cluster Is Created

Perform the following steps:

1. Run `netconfig` to configure the network interfaces that you want to use on the private subnet. You must configure a network interface for each member that you want to use on the private subnet.

Assign a host name and use an IP address reserved for private subnets. The example uses a host name that reflects the host name of the interface on the public network. The host portion of the IP address is the same as that used on the public network to make it easy to identify.

```
interface type: ee2
host name for this interface: polishhamx.zk3.dec.com
IP address for this interface: 172.16.112.176
```

2. When prompted by `netconfig`, restart the network services on each cluster member on which you configured a network interface to use on the private subnet.
3. Edit the clusterwide `/etc/hosts` file and add the cluster alias for the private subnet. Specify the IP address of the host, the fully qualified host name including the domain name, and the unqualified host name for this host:

```
172.16.112.209 delix.zk3.dec.com delix
```

4. Join the cluster alias you want to use for the cluster on the private subnet. To do this, execute the following command on each member that you want to join the alias:  

```
cluamgr -a alias=delix.zk3.dec.com,join
```
5. On each member that you want to join the alias, edit the `/etc/clu_alias.config` file and add an entry for the cluster alias on the private subnet:

```
/usr/sbin/cluamgr -a selw=3,selp=1,join,alias=delix.zk3.dec.com
```

The `/etc/clu_alias.config` file is a context-dependent symbolic link (CDSL) pointing to member-specific cluster alias configuration files. (The `clu_create` and `clu_add_member` commands add the

default cluster entry to the file.) See `cluamgr(8)` for a description of the `selw` and `selp` options.

## Monitoring the Private Subnet with CAA

You can use the cluster application availability (CAA) subsystem to create and monitor a network resource that includes the private subnet. The `caa_profile -create resource_name -t network -s subnet_address` command creates a resource profile for the indicated resource. In particular, the `-s` option specifies the subnet address of the network resource. CAA then indicates whether the network resource exceeds the failure interval or failure threshold.

The following example creates a profile for a network resource. The network is the bitwise AND of the IP address and the netmask. If you assume an IP address of 172.16.112.176 and a netmask of 255.255.255.0, the subnet will be 172.16.112.0:

```
/usr/sbin/caa_profile -create network1 -t network -s 172.16.112.0
/usr/sbin/caa_register network1
/usr/sbin/caa_profile -print network1
```

```
NAME=network1
TYPE=network
DESCRIPTION=network1
FAILURE_INTERVAL=0
FAILURE_THRESHOLD=0
SUBNET=172.16.112.0
```

You can then use the `caa_stat` command either interactively or from a script to test the status of the network resource:

```
/usr/bin/caa_stat network1
```

## Verifying Success

After you apply this Best Practice for configuring a dual-alias cluster on public and private subnets, you can verify whether it was successful.

The cluster systems share the `/etc/hosts` file and should be able to access each other from both the private subnet or public subnet. Use `rlogin` or a simple X application such as `xclock` or `xterm` to verify this.

The alternate cluster alias and all of its members should be inaccessible from any other system on the public subnet. Use `rlogin` or a simple X application such as `xclock` or `xterm` to verify the inaccessibility of the systems.

At the same time, the default cluster alias (`deli` in the example) and all of its members should be accessible from the public subnet.

Use the alternate cluster alias and its members' names when installing any application software that can be configured to run on a specific host name or network interface.

If the Best Practice was not successful, see *Troubleshooting* for information about identifying and solving problems.

## Troubleshooting

If you determine that this Best Practice was not successful, as described in *Verifying Success*, use the following steps to identify and solve problems:

1. Verify that the `/etc/hosts` file contains entries for the cluster alias and member host names. For example,

```
172.16.112.176 polishhamx.zk3.dec.com polishhamx
172.16.112.238 provolonex.zk3.dec.com provolonex
172.16.112.209 delix.zk3.dec.com delix
```

2. Use the `netstat -i` command to make sure the network interfaces for the private subnet are defined correctly.
3. Use the `cluamgr -s alias` command to make sure the cluster alias for the private subnet is defined correctly.

## Comments and Questions

We value your comments and questions on the information in this document. Please mail your comments to us at this address:

`best_practices@zk3.dec.com`

## Legal Notice

Compaq, the Compaq logo, Tru64, and TruCluster are trademarks of Compaq Information Technologies Group, L.P. in the U.S. and/or other countries. UNIX is a trademark of The Open Group in the U.S. and/or other countries. All other product names mentioned herein may be trademarks of their respective companies.

Confidential computer software. Valid license from Compaq 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.

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