

Introduction to Oracle® Solaris 11 Networking

Copyright © 2011, 2012, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. UNIX est une marque déposée concédé sous licence par X/Open Company, Ltd.

Contents

Preface	5
1 Networking on Oracle Solaris 11 Systems	7
Network Configuration Profiles	7
Reactive Network Configuration	8
Fixed Network Configuration	8
System-Generated Profiles	8
Network Stack in Oracle Solaris	9
Previous Oracle Solaris Releases	9
Oracle Solaris 11 Implementation	11
Network Devices and Datalink Names	15
Default Generic Link Names	15
Assignment of Generic Names to Datalinks	16
Customizing How Generic Link Names Are Assigned By the Operating System	17
Link Names on Upgraded Systems	18
Index	21

Preface

Welcome to Introduction to Oracle Solaris 11 Networking. This book is part of the series *Establishing An Oracle Solaris 11.1 Network* that cover basic topics and procedures to configure Oracle Solaris networks. This book assumes that you have already installed Oracle Solaris. You should be ready to configure your network or ready to configure any networking software that is required on your network.

Who Should Use This Book

This book is intended for anyone responsible for administering systems that run Oracle Solaris, which are configured in a network. To use this book, you should have at least two years of UNIX system administration experience. Attending UNIX system administration training courses might be helpful.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-1 Typographic Conventions

Typeface	Description	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:

TABLE P-1 Typographic Conventions (Continued)

Typeface	Description	Example
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows UNIX system prompts and superuser prompts for shells that are included in the Oracle Solaris OS. In command examples, the shell prompt indicates whether the command should be executed by a regular user or a user with privileges.

TABLE P-2 Shell Prompts

Shell	Prompt
Bash shell, Korn shell, and Bourne shell	\$
Bash shell, Korn shell, and Bourne shell for superuser	#
C shell	machine_name%
C shell for superuser	machine_name#

Networking on Oracle Solaris 11 Systems

This book is an introduction to networking in Oracle Solaris. It describes features that are a foundation for the manner in which you configure the network on systems that run the Oracle Solaris 11 Operating System.

This book does not cover all the network topics necessary for a basic understanding of networking, such as the TCP/IP stack, IP addresses, CIDR notations, subnets, and other basic concepts. For information about these networking concepts, refer to any introductory book on networking, including Oracle Solaris administration guides such as [Chapter 1, “Oracle Solaris TCP/IP Protocol Suite \(Overview\)”](#), in *System Administration Guide: IP Services*.

For information about receiving training on Oracle Solaris networking, go to the [Oracle Solaris 11 Training page \(http://www.oracle.com/us/education/selectcountry-new-079003.html\)](http://www.oracle.com/us/education/selectcountry-new-079003.html).

This book discusses the following topics:

- “Network Configuration Profiles” on page 7
- “Network Stack in Oracle Solaris” on page 9
- “Network Devices and Datalink Names” on page 15

Network Configuration Profiles

In Oracle Solaris 11, network configuration profiles (NCPs) manage the system's network configuration. Only one NCP can be active on a system at a time. Oracle Solaris 11 supports two types of NCPs: *reactive* and *fixed*. Depending on its type, the active NCP determines whether the system's network configuration is reactive or fixed.

Reactive Network Configuration

In reactive network configuration, a network daemon monitors the system's network configuration. If networking conditions change for the system, then its network configuration changes for the new conditions. For example, suppose that a system with multiple network interface cards (NICs) is connected to two networks. In one network, no Dynamic Host Configuration Protocol (DHCP) server is available. In this situation, a reactive network configuration can be implemented on the system. Two reactive NCPs are defined. One NCP manages the system's connection to the network that has DHCP support. The other NCP manages the connections when DHCP support is absent. The daemon automatically activates one or the other NCP depending on the networking conditions in which the system is currently operating. With reactive configuration, the system adjusts to the two network setups automatically without requiring manual reconfiguration.

You can create multiple reactive NCPs on a system to match varying network setups for the system. However, only one NCP together with a corresponding Location profile can be active at a time. You set up policies to determine which reactive NCP is activated for a given condition.

For more information about NCPs, see [Connecting Systems Using Reactive Network Configuration in Oracle Solaris 11.1](#).

Fixed Network Configuration

Fixed network configuration is the opposite of reactive network configuration. The network daemon instantiates a specific network configuration on the system, but does not automatically adjust that configuration to varying conditions. To implement fixed networking, a fixed profile must be activated on the system. A system can have multiple reactive profiles for reactive networking. However, only one fixed profile exists on a system.

Fixed network configuration does not refer only to using static IP addresses. On a system on which reactive network configuration is implemented, you can create a reactive profile where the IP interfaces are assigned static IP addresses. The profile is used, for example, when no DHCP service is available.

Thus, reactive and fixed network configurations must be understood in terms of the ability of a system to automatically adjust to changes in its networking conditions rather than simply on whether fixed or static IP addresses are used.

System-Generated Profiles

On a system that is running Oracle Solaris 11.1, two NCPs are automatically created by the system: `Automatic` for reactive network configuration and `DefaultFixed` for fixed network configuration. The NCP that you enable during installation becomes the active NCP for the

system. In turn, the active NCP determines the type of network configuration. For information about installing Oracle Solaris 11.1, see “[Networking Configuration With Text Installer](#)” in *Installing Oracle Solaris 11.1 Systems*.

Aside from the `Automatic` NCP, you can create other reactive profiles for reactive networking. However, you cannot have any other fixed profiles on the system except `DefaultFixed`.

After you have installed Oracle Solaris, you can switch between reactive and fixed networking by using the `netadm` command.

```
# netadm enable -p ncp ncp-name
```

where *ncp-name* refers to the NCP that you want to make active.

Network Stack in Oracle Solaris

Network interfaces provide the connection between the system and the network. These interfaces are configured over datalinks, which in turn correspond to instances of hardware devices on the system. Network hardware devices are also called *network interface cards (NICs)* or *network adapters*. NICs can be built in and already present on the system when the system is purchased. However, you can also purchase separate NICs to add to the system. Certain NICs have only a single interface that resides on the card. Other brands might have multiple interfaces that you can configure to perform network operations.

Previous Oracle Solaris Releases

In previous Oracle Solaris implementations of the network stack, interfaces and links on the software layer build on the devices in the hardware layer. More specifically, a hardware device instance in the hardware layer has a corresponding link on the datalink layer and a configured interface on the interface layer. This one-to-one relationship among the network device, its datalink, and the IP interface is illustrated in the figure that follows.

FIGURE 1-1 Network Stack Showing Network Devices, Links, and Interfaces — Previous Oracle Solaris Implementation

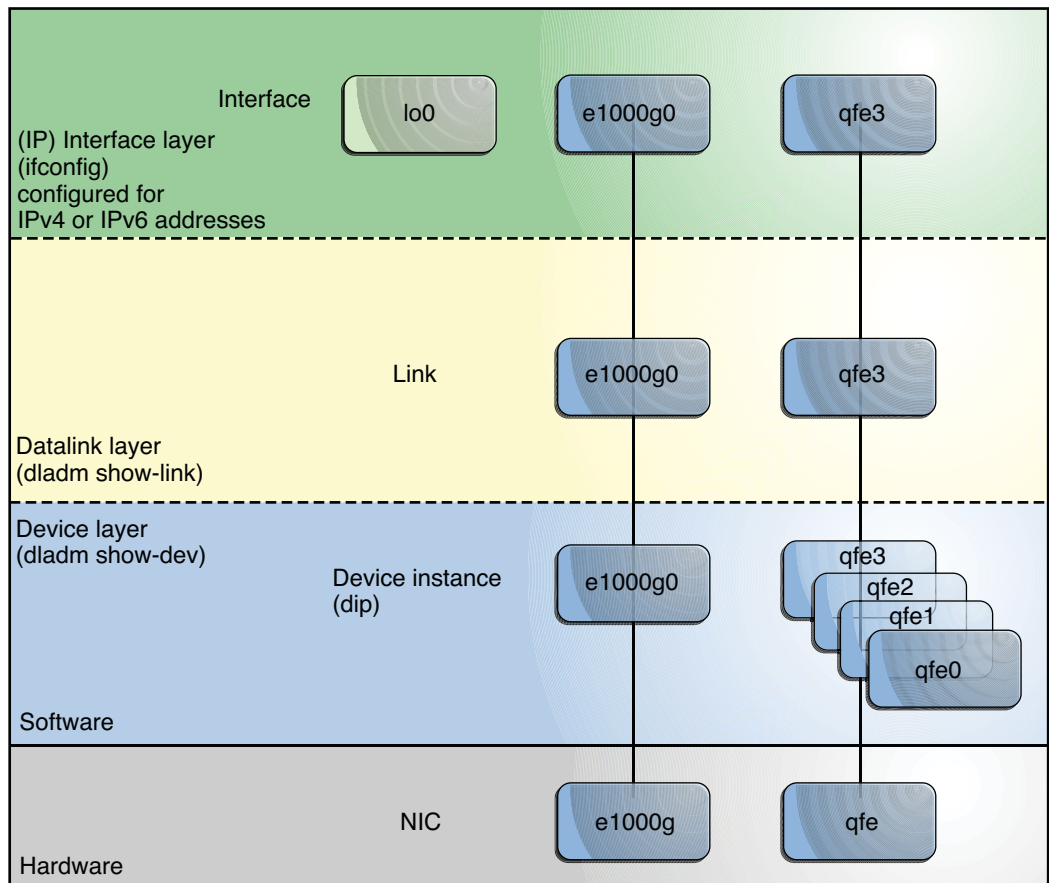


Figure 1-1 shows two NICs on the hardware layer: `e1000` with a single device instance `e1000g0`, and `qfe` with multiple device instances, `qfe0` to `qfe3`. The devices `qfe0` through `qfe2` are not used. Devices `e1000g` and `qfe3` are used and have corresponding links `e1000g` and `qfe3` on the datalink layer. In the figure, the IP interfaces are likewise named after their respective underlying hardware, `e1000g` and `qfe3`. These interfaces can be configured with IPv4 or IPv6 addresses to host both types of network traffic. Note also the presence of the loopback interface `lo0` on the interface layer. This interface is used to test, for example, that the IP stack is functioning properly.

Different administrative commands are used at each layer of the stack. For example, the `dladm show-dev` command lists hardware devices that are installed on the system. The `dladm show-link` command lists information about links on the datalink layer. The `ifconfig` command shows the IP interface configuration on the interface layer.

In this model, a one-to-one relationship exists that binds the device, the datalink, and the interface. This relationship means that network configuration is dependent on hardware configuration and network topology. Interfaces must be reconfigured if changes are implemented in the hardware layer, such as replacing the NIC or changing the network topology.

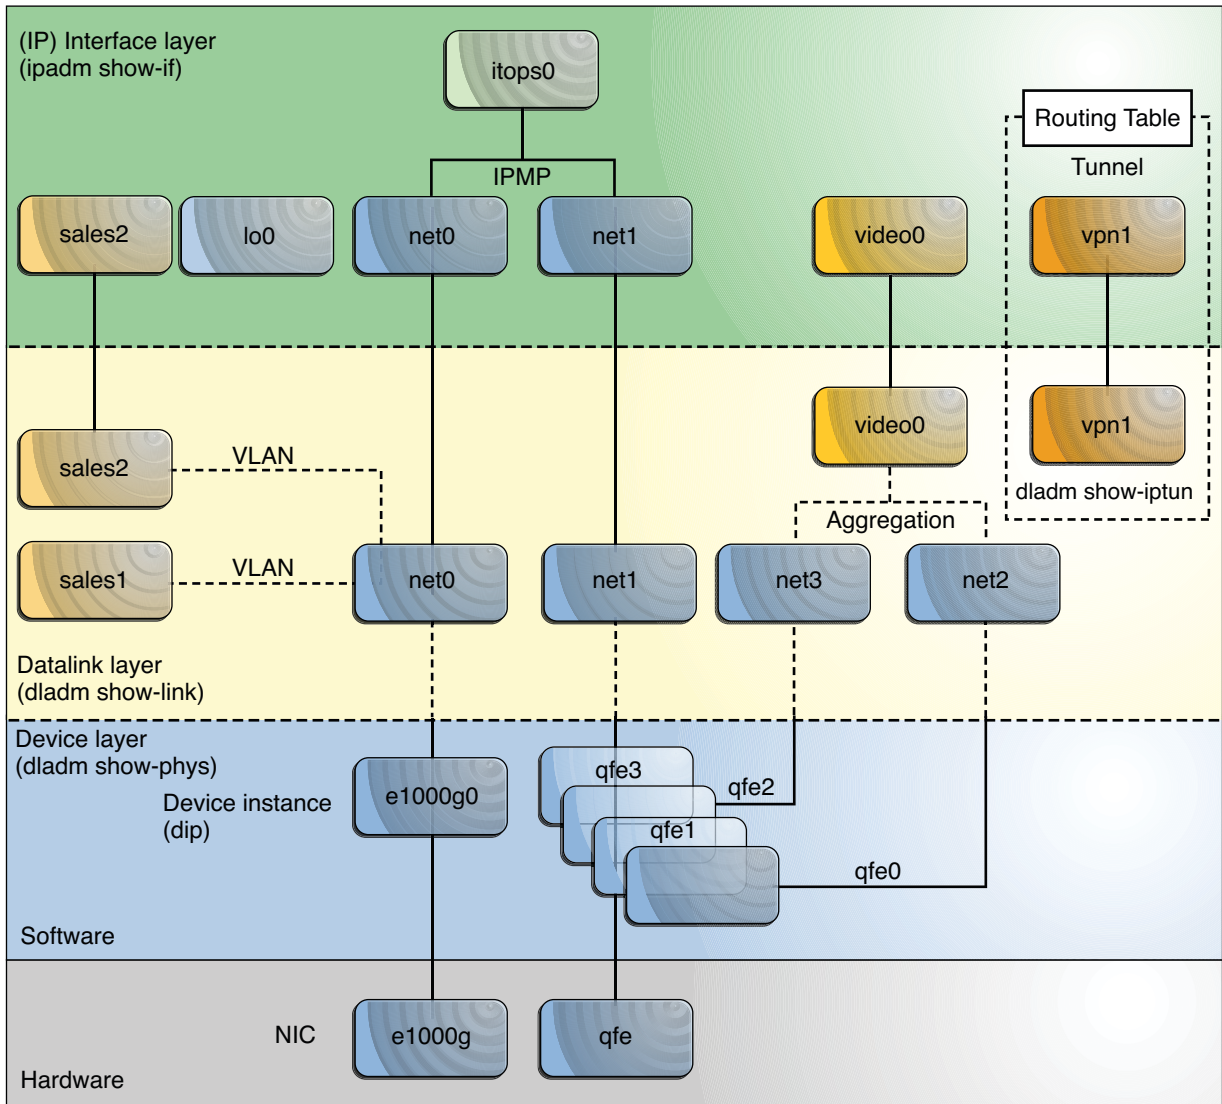
Oracle Solaris 11 Implementation

In Oracle Solaris 11, the one-to-one relationship between the hardware, datalink, and interface layers remains. However, the software layer is decoupled from the hardware layer. With this separation, network configuration on the software level is no longer bound to the chipset or the network topology in the hardware layer. This implementation makes network administration more flexible in the following ways:

- The network configuration is insulated from any changes that might occur in the hardware layer. Link and interface configurations are preserved even if the underlying hardware is removed. These same configurations can then be reapplied to any replacement NIC, provided that the two NICs are of the same type.
- The separation of the network configuration from the network hardware configuration also allows the use of customized link names in the datalink layer.
- With the abstraction of the datalink layer, multiple networking abstractions or configurations such as VLANs, VNICs, physical devices, link aggregations, and IP tunnels are unified into a common administrative entity, which is the datalink.

The following figure shows the interrelationship among devices, link types, and their corresponding interfaces.

FIGURE 1-2 Network Stack Showing Network Devices, Links, and Interfaces — Oracle Solaris 11 Implementation



Note – In the figure, the datalinks are named according to specific functions that they perform on the system, such as `video0` or `sales2`. The figure highlights the flexibility with which you can name the datalinks. However, using the default neutral names such as `net0` as supplied by the OS is sufficient and preferable. Datalink names are discussed in “[Network Devices and Datalink Names](#)” on page 15.

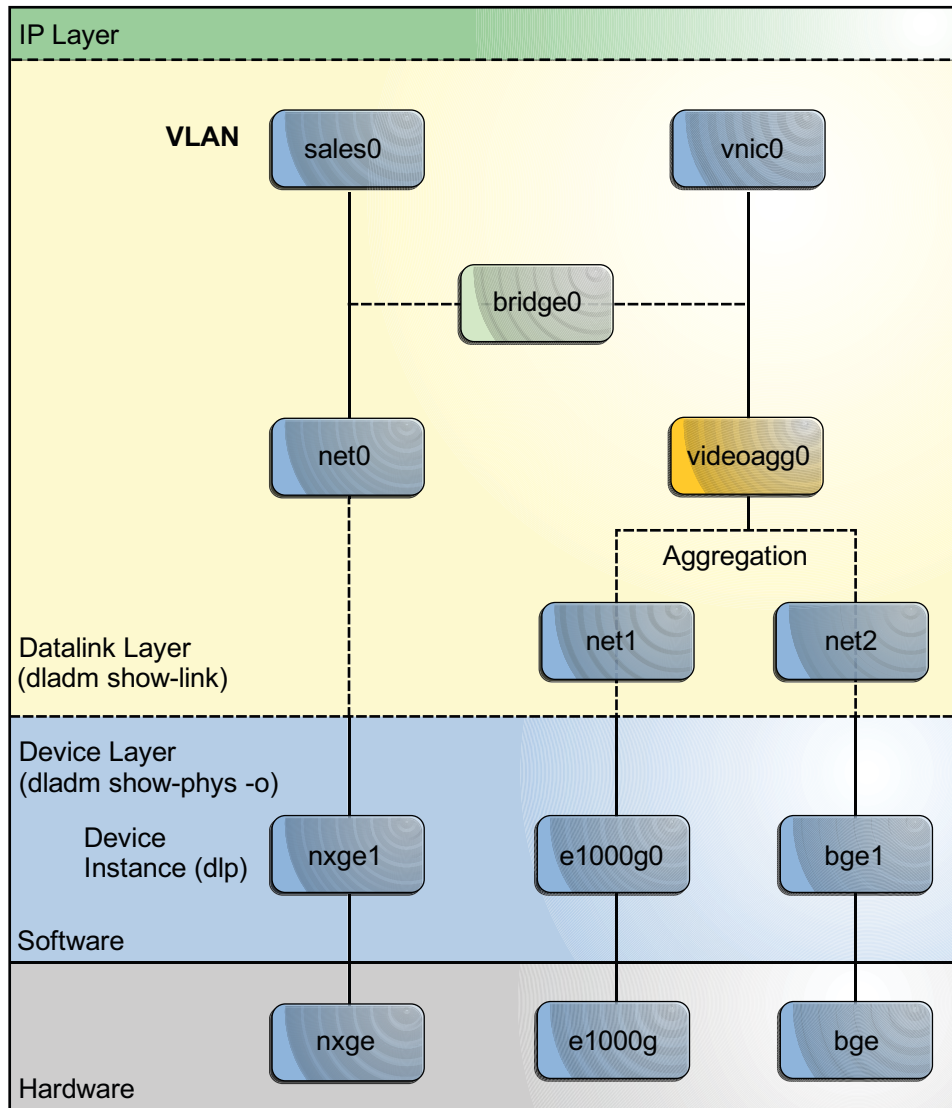
Figure 1–2 also provides a sample of how administratively chosen names can be used in the network setup.

- Virtual local area networks (VLANs) are configured on the `net0` link. These VLANs, in turn, are also assigned customized names, such as `sales1` and `sales2`. The VLAN `sales2`'s IP interface is plumbed and operational.
- The datalinks of the device instances `qfe0` and `qfe2` are automatically named by the OS. These two datalinks are aggregated to host video feed. The aggregation's name can be customized. In the figure, the aggregation is named `video0`.
- Two interfaces (`net0` and `net1`) with different underlying hardware (`e1000g` and `qfe`) are grouped together as an IP multipathing (IPMP) group called `itops0`.
- Two interfaces have no underlying devices: The tunnel `vpn1` is configured for VPN connections, and `lo0` exists for IP loopback operations.

All of the link and interface configurations in this figure are independent of the configurations in the underlying hardware. For example, if the `qfe` card is replaced, the `video0` interface configuration for video traffic remains and can later be applied to a replacement NIC.

On the datalink layer of the same network stack implementation, you can configure bridges as shown in the following figure. Two interfaces, `net0` and `videoagg0`, are configured as a bridge, `bridge0`. Packets that are received on one interface are forwarded to the other interface. After bridge configuration, both interfaces can still be used to configure VLANs and IP interfaces.

FIGURE 1-3 Bridges in the Network Stack for Oracle Solaris 11



Bridges and bridge configuration are discussed in Chapter 4, “Administering Bridged Networks (Tasks)” in *Managing Oracle Solaris 11.1 Network Performance*.

Network Devices and Datalink Names

From an administrative perspective, administrators create IP interfaces on top of datalinks. The datalink represents a link object in the second layer of the Open Systems Interconnection (OSI) model. The *physical link* is directly associated with a device and possesses a device name. The device name is essentially the device instance name, which contains the driver name and the device instance number. The instance number can have a value from zero to $n-1$, depending on how many NICs use that driver on the system.

For example, consider a Gigabit Ethernet card, which is often used as the primary NIC on both host systems and server systems. Some typical driver names for this NIC are `bge` and `e1000g`. When used as the primary NIC, the Gigabit Ethernet interface has a device name such as `bge0` or `e1000g0`. Other driver names are `nge`, `nxge`, and so on.

In this Oracle Solaris release, the device instance name continues to depend on the underlying hardware. However, as a consequence of the separation between hardware and software layers, datalinks on top of these devices are not similarly bound. Thus these datalinks can be given names other than the device names on which they are configured.

Default Generic Link Names

In Oracle Solaris 11, datalinks are automatically provided with generic names by default. This name assignment uses the `net#` naming convention, where `#` is the instance number. This instance number increments for each device, for example, `net0`, `net1`, `net2`, and so on.

Generic or flexible link names provide advantages in network configuration, as shown in the following examples:

- Within a single system, dynamic reconfiguration (DR) becomes easier. The network configuration for a given NIC can be inherited by a different NIC replacement.
- Zone migration becomes less complicated with regards to network setup. The zone in the migrated system preserves its network configuration if the destination system's link shares the same name with the link that has been assigned to the zone prior to migration. Thus, no additional network configuration on the zone is required after the migration.
- The generic naming scheme helps with the network configuration that is specified in the System Configuration (SC) manifest. The primary network datalink is generally named `net0` for all systems. Thus, a generic SC manifest can be used for multiple systems that specify a configuration for `net0`.
- Datalink administration also becomes flexible. You can further customize the name of datalinks, for example, to reflect a specific function that the datalink serves, as shown in [Figure 1–2](#).

The following table illustrates the new correspondence between the hardware (NIC), the device instance, the link name, and the interface over the link. The names of the datalinks are automatically provided by the OS.

Hardware (NIC)	Device Instance	Link's Assigned Name	IP Interface
e1000g	e1000g0	net0	net0
qfe	qfe1	net1	net1

As the table indicates, while the device instance name remains hardware-based, the datalinks are renamed by the OS after it is installed.

To display the mapping between the datalinks with their generic names and the corresponding device instances, you use the `dladm show-phys` subcommand. For example:

```
# dladm show-phys
LINK      MEDIA      STATE      SPEED      DUPLEX     DEVICE
net2      Ethernet   up         1000      full       bge2
net0      Ethernet   up         1000      full       e1000g0
net3      Ethernet   up         1000      full       nge3
net1      Ethernet   up         1000      full       e1000g1
```

Assignment of Generic Names to Datalinks

In Oracle Solaris, the OS provides the generic names for all the datalinks based on specific criteria. All devices share the same prefix `net`. However, the instance numbers are assigned based on the following:

- Physical network devices are ordered according to media type, where certain types have priority over others. The media types are ordered in descending priority as follows:
 1. Ethernet
 2. IP over IB (Infiniband devices)
 3. Ethernet over IB
 4. WiFi
- After devices are grouped and sorted according to media types, these devices are further ordered based on their physical locations, where on-board devices are favored over peripheral devices.
- Devices that have higher priority based on their media type and location are assigned lower instance numbers.

Based on the criteria, Ethernet devices on a lower motherboard or IO board, host bridge, PCIe root complex, bus, device, and function are ranked ahead of the other devices.

To display the correspondences of link names, devices, and locations, use the `dladm show-phys` command, as follows:

```
# dladm show-phys -L
LINK      DEVICE      LOCATION
net0      e1000g0    MB
```


net1	e1000g1	MB
net2	e1000g2	MB
net3	e1000g3	MB
net4	ibp0	MB/RISER0/PCIE0/PORT1
net5	ibp1	MB/RISER0/PCIE0/PORT2
net6	eoib2	MB/RISER0/PCIE0/PORT1/cloud-nm2gw-2/1A-ETH-2
net7	eoib4	MB/RISER0/PCIE0/PORT2/cloud-nm2gw-2/1A-ETH-2

Customizing How Generic Link Names Are Assigned By the Operating System

Oracle Solaris uses the prefix `net` when assigning link names. However, any custom prefix can be used instead, such as `eth`. If you prefer, you can also disable the automatic assignment of neutral link names.



Caution – You must customize how generic link names are automatically assigned **before** you install Oracle Solaris. After installation, you cannot customize the default link names without removing existing configurations.

To disable automatic link naming, or to customize the prefix of link names, set the following property in the System Configuration (SC) manifests. The SC manifests are used by the Automated Installer (AI) feature of Oracle Solaris.

```
<service name="network/datalink-management"
  version="1" type="service">
  <instance name="default enabled="true">
    <property_group name='linkname-policy'
      type='application'>
      <propval name='phys-prefix' type='astring'
        value='net' />
    </property_group>
  </instance>
</service>
```

By default, the value for `phys-prefix` is set to `net`, as shown in bold.

- To disable automatic naming, remove any value that is set for `phys-prefix`. If you disable automatic naming, then datalink names are based on their associated hardware drivers, such as `bge0`, `e1000g0`, and so on.
- To use a prefix other than `net`, specify a new prefix as the value of `phys-prefix`, such as `eth`.

If the value that is provided to `phys-prefix` is invalid, then that value is ignored. The datalinks are then named according to their associated hardware drivers, such as `bge0`, `e1000g0`, and so on. For rules about valid link names, see [“Rules for Valid Link Names” on page 19](#).

Link Names on Upgraded Systems

On systems where the Oracle Solaris 11 release is freshly installed, datalinks are automatically named `net0` through `netN-1`, where *N* represents the total number of network devices.

On the contrary, if you upgrade from Oracle Solaris 11 Express, the datalinks retain their names that were established prior to the upgrade. These names are either be the default hardware-based names or customized names that the administrator assigned to the datalinks before the upgrade. Further, on these upgraded systems, new network devices that are subsequently added also retain the default hardware-based names rather than receive generic names. This behavior for upgraded systems ensures that no generic names assigned by the OS become mixed with other hardware-based names or customized names assigned by the administrator before the upgrade.

On any system with Oracle Solaris 11, both hardware-based names as well as OS-supplied link names can be replaced by other names that you prefer to use. Typically, the default link names that are assigned by the OS suffice for creating the system's network configuration. However, to change link names, note the important considerations discussed in the following sections.

Replacing Hardware-Based Link Names

If your system's links have hardware-based names, rename these links with at least generic names. If you retain the hardware-based names, confusion might arise later when these physical devices are removed or replaced.

For example, you retain the link name `bge0` that is associated with the device `bge0`. All link configurations are performed by referring to the link name. Later, you might replace the NIC `bge` with the NIC `e1000g`. To reapply the former device's link configuration to the new NIC `e1000g0`, you would need to reassign the link name `bge0` to `e1000g0`. The combination of a hardware-based link name `bge0` with a different associated NIC `e1000g0` can cause confusion. By using names that are not hardware-based, you can better distinguish the links from the associated devices.

Caution About Changing Link Names

Although replacing hardware-based link names is a best practice, you must still plan carefully before you rename links. Changing a device's link name does not automatically propagate the new name to all existing associated configurations. The following examples illustrate the risks when you change link names:

- Some rules in an IP Filter configuration apply to specific links. When you change a link's name, the filter rules continue to refer to the link's original name. Consequently, these rules no longer behave as expected after you rename the link. You need to adjust the filter rules to apply to the link by using the new link name.

- Consider the possibility of exporting network configuration information. As previously explained, by using the default `net#` names provided by the OS, you can migrate zones and easily export network configuration to another system. If the target system's network devices are named with generic names such as `net0`, `net1`, and so on, then the zone simply inherits the network configuration of the datalink whose name matches the datalink assigned to the zone.

Thus, as a general rule, do not rename datalinks randomly. When renaming datalinks, ensure that all of the link's associated configurations continue to apply after the link name is changed. Some of the configurations that might be affected by renaming links are as follows:

- IP Filter rules
- IP configurations that are specified in configuration files such as `/etc/dhcp.*`
- Oracle Solaris 11 zones
- autopush configuration

Note – No changes are required in the autopush configuration when you rename links. However, you must be aware of how the configuration works with the per-link autopush property after the link has been renamed. For more information, see “[Setting the STREAMS Module on Datalinks](#)” in *Connecting Systems Using Fixed Network Configuration in Oracle Solaris 11.1*.

Rules for Valid Link Names

When you assign link names, observe the following rules:

- Link names must consist of a string and a *physical point of attachment (PPA)* number.
- The link name must abide by the following constraints:
 - Names ideally consist of between 3 to 8 characters. However, names can have a maximum of 16 characters.
 - Valid characters for names are alphanumeric (a–z, 0–9) and the underscore (_).



Caution – Do not use upper case letters on link names.

- Each datalink must have only one link name at one time.
- Each datalink must have a unique link name within the system.

Note – As an added restriction, you cannot use `lo0` as a flexible link name. This name is reserved to identify the IP loopback interface.

The function of the link within your network setup can be a useful reference when you assign link names. For example, `netmgmt0` can be a link that is dedicated to network management. `Upstream2` can be the link that connects to the ISP. As a general rule to avoid confusion, do *not* assign names of known devices to your links.

Index

A

Automated Installer (AI), 17
Automatic NCP, 8–9

B

bridges, 14

C

customized names, 15–20

D

datalinks

- link names, 15–20
- naming conventions, 15–20
- rules for creating link names, 19–20
- tunnel links, 11
- VLANs, 11
- VNICs, 11

DefaultFixed NCP, 8–9

device instance, 9, 11

device locations, 16

dladm command, 8

- displaying device locations, 16
- show-dev, 9–11
- show-link, 11
- show-phys, 11, 16

dynamic reconfiguration (DR), 15

E

Ethernet devices, 16

F

fixed network configuration, 8

I

ifconfig command, 9–11

Infiniband, 16

ipadm command, 8

- show-if, 11

L

link names, 15–20

- automatic naming, 17

- changing, 18–19

- customizing, 17

- dynamic reconfiguration (DR), 15

- generic, 15–16, 16–17

- loopback interface, 19

- on upgraded systems, 18–20

- replacing hardware-based names, 9, 18

- System Configuration (SC) manifest, 15

- zone migration, 15

loopback interface, 19

N

NCPs, *See* network configuration profiles

netadm command, 8

netcfg command, 8

network configuration

- fixed, 8
- reactive, 8

network configuration profiles, 7–9

- Automatic NCP, 8–9
- DefaultFixed NCP, 8–9
- switching, 8–9
- system-generated, 8–9

network devices, 16

- Ethernet, 16
- Infiniband, 16

network interface cards, 8, 9–14

network stack, 9–14, 11

- datalinks, 9–11
- datalinks display, 9, 11
- Oracle Solaris 11, 11–14
 - customized names, 11
 - datalinks, 11–14
 - hardware and software layer, 11–14
 - IP interfaces, 11–14
 - listing IP interfaces, 11
 - showing physical links and devices, 11
- previous Oracle Solaris releases, 9–11
 - hardware and software layer, 9–11
 - IP interfaces, 9–11
 - listing IP interfaces, 9–11
 - showing devices, 9–11

networking configuration commands

- dladm, 8
- ipadm, 8
- netadm, 8
- netcfg, 8

P

physical point of attachments, 19–20

profile-managed network configuration, 7–9

R

reactive network configuration, 8

S

System Configuration (SC) manifest, 15

T

tunnels, 11

V

virtual local area networks (VLANs), 11

virtual network interface cards (VNICs), 11

W

WiFi, 16

Z

zone migration, 15