

Sun Ethernet Fabric Operating System

LLA Administration Guide



Part No.: E41876-01
July 2013

Copyright © 2013, 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 software 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.

Copyright © 2013, Oracle et/ou ses affiliés. Tous droits réservés.

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 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.

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.

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. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.



Contents

Using This Documentation 1

- Product Notes 1
- Related Documentation 2
- Acronyms and Abbreviations 2
- CLI Command Modes 2
- Feedback 3
- Support and Accessibility 3

LLA Overview 5

- Protocol Description 5
- LLA Port-Types 6
- LLA Port-Channels 6
- Supported CLI Commands 7

Configuring LLA 9

- LLA Example Topology 9
- ▼ Configure LLA 9

Using This Documentation

LLA is a L2 feature providing the aggregation of links from a server connected to two physical switches resulting in increased redundancy and capacity. When the port connecting to the server to one of the switches fails or is shutdown, LLA supports fail-over by redirecting the traffic over the ISL. This document describes the basic configuration tasks of SEFOS LLA.

- “Product Notes” on page 1
- “Related Documentation” on page 2
- “Acronyms and Abbreviations” on page 2
- “CLI Command Modes” on page 2
- “Feedback” on page 3
- “Support and Accessibility” on page 3

Product Notes

For late-breaking information and known issues about the following products, refer to the product notes at:

Oracle Switch ES1-24:

<http://www.oracle.com/goto/ES1-24/docs>

Sun Network 10GbE Switch 72p:

<http://www.oracle.com/goto/SN-10GbE-72p/docs>

Sun Blade 6000 Ethernet Switched NEM 24p 10GbE:

<http://www.oracle.com/goto/SB6K-24p-10GbE/docs>

Related Documentation

Documentation	Links
All Oracle products	http://oracle.com/documentation
Oracle Switch ES1-24	http://www.oracle.com/goto/ES1-24/docs
Sun Network 10GbE Switch 72p	http://www.oracle.com/goto/SN-10GbE-72p/docs
Sun Blade 6000 Ethernet Switched NEM 24p 10GbE	http://www.oracle.com/goto/SB6K-24p-10GbE/docs
Sun Blade 6000 modular system	http://www.oracle.com/pls/topic/lookup?ctx=sb6000
Oracle Integrated Lights Out Manager (Oracle ILOM) 3.0	http://www.oracle.com/pls/topic/lookup?ctx=ilom30

For detailed information about the commands and options described in this document, refer to the *Sun Ethernet Fabric Operating System CLI Base Reference Manual*.

Acronyms and Abbreviations

Acronym or Abbreviation	Explanation
LLA	Logical link aggregation
ISL	Inter-switch link
SEFOS	Sun Ethernet Fabric Operating System

CLI Command Modes

The following table lists the configuration modes used in this document with their access and exit methods.

Command Mode	Access Method	Prompt	Exit Method
User EXEC	Access SEFOS from Oracle ILOM with read-only rights (privilege level 1).	SEFOS>	Use the <code>logout</code> or <code>exit</code> command to return to the Oracle ILOM prompt.
Privileged EXEC	Access SEFOS from Oracle ILOM with full administrative rights (privilege level 15).	SEFOS#	Use the <code>logout</code> or <code>exit</code> command to return to the Oracle ILOM prompt.
Global Configuration	From User EXEC mode, use the <code>enable</code> command.	SEFOS(config)#	Use the <code>end</code> command to return to Privileged EXEC mode.
Interface Configuration	From Global Configuration mode, use the <code>interface interface-type interface-id</code> command.	SEFOS(config-if)#	Use the <code>exit</code> command to return to Global Configuration mode, or use the <code>end</code> command to return to Privileged EXEC mode.
LLA Configuration	From Global Configuration mode, use the <code>lla</code> command.	SEFOS(config-lla)#	Use the <code>exit</code> command to return to Global Configuration mode, or use the <code>end</code> command to return to Privileged EXEC mode.

Feedback

Provide feedback on this documentation at:

<http://www.oracle.com/goto/docfeedback>

Support and Accessibility

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.

LLA Overview

This section describes the LA protocol.

- [“Protocol Description” on page 5](#)
- [“LLA Port-Types” on page 6](#)
- [“LLA Port-Channels” on page 6](#)
- [“Supported CLI Commands” on page 7](#)

Protocol Description

Standard link aggregation enables you to directly connect two systems through multiple ports. These ports are viewed as one logical link, and may be simultaneously in active mode. This situation provides higher capacity between the two systems and supports automatic fail over if any of the physical ports becomes unavailable. Typically, you can use link aggregation between tiers for higher capacity. To provide redundancy, multiple up-links are deployed between tiers. However, because of the spanning tree algorithm, only one of these links is active.

LLA on the other hand, enables link aggregation links from servers to terminate at two separate switches. This situation provides an increase in capacity and redundancy, supporting an active-active deployment model. Fail-over can occur rapidly, as there are no L2 or L3 protocol delays. On the server, the links are configured as Linux bonds or Oracle Solaris link aggregations.

In LLA, two switches are connected together directly through an ISL, a special trunk group between them. The LLA ports are connected to the server and are members of the LLA port-channels. The ISL port-channel is required to be in forwarding state on both the switches at all times. This is done by setting the spanning tree cost on the ISL port-channel lower than any other port in the system, and by choosing the uplink costs such that the sum of the ISL cost and the preferred uplink cost is less than that of other uplinks. Both switches maintain their identity as separate L2/L3 switching entities. The rest of the network sees them as two separate L2 switches with different addresses.

When designing a deployment with LLA, carefully examine the expected traffic patterns. The ISL port-channel should be provisioned with an adequate bandwidth. Traffic from the uplinks may always be distributed between the local link and the path crossing the ISL. Likewise, depending on the server side configuration, traffic may always be using both paths in the reverse direction.

With LLA enabled, the switches must share state and maintain a common mapping of ports. This common mapping is initialized when SEFOS starts up. For this reason, after LLA is configured, the configuration must be saved and SEFOS must be restarted.

LLA Port-Types

There are three types of LLA ports:

- **LLA ports.** Ports connected to the hosts and part of the LLA port-channel.
- **ISL ports.** Ports used to connect the two switches and part of the ISL port-channel. The ISL ports have to be consecutive, otherwise an error is flagged and LLA is disabled on subsequent start-up.
- **Standard ports.** Any other ports. Standard ports may be part of regular port-channels.

LLA Port-Channels

The port-channels could be any of the following three types:

- **LLA port-channel.** Configured on both systems with the same port-channel key and port number. The LLA port-channel key range is from 101 to 124. Only one LLA port per LLA port-channel is allowed. The range of LLA port-channels from 125 to 130 is reserved for internal usage. New LLA port-channels are allowed at run-time. However, changes to existing LLA ports and LLA port-channel is not supported.
- **ISL port-channel.** Configured on both systems and member ports connecting the two switches. The ISL port-channel key is fixed as 100. When multiple VLANs are configured, ISL port-channel should be in trunk mode. The ISL port-channel must be a member of every VLAN that the LLA port-channels are in. Changes to the ISL port-channel will require SEFOS to be restarted after saving the configuration.

- **Regular port-channel.** Any standard ports not configured as either LLA or ISL ports can be member of regular port-channels. Up to 12 regular port-channels can be created, with the port-channel keys ranging from 131 to 142.

Supported CLI Commands

You can use the following CLI commands to configure LLA:

- `set lla enable`
- `set lla disable`
- `lla`
- `isl port-channel key`
- `lla port-channel key`
- `no isl port-channel key`
- `no lla port-channel key`
- `role primary`
- `secondary primary`
- `no role`
- `show lla role`
- `show lla ports`
- `show lla status`


Refer to the *Sun Ethernet Fabric Operating System CLI Base Reference Manual* for the complete set of commands for configuring LLA.

Configuring LLA

These sections describe how to configure LLA on a switch.

- “LLA Example Topology” on page 9
- “Configure LLA” on page 9

LLA Example Topology



▼ Configure LLA

To configure LLA on a switch, you must first configure the LLA and ISL port-channels.

Connect the host to one port on each of the switches. The connections must be to the same port number on both switches. Form an aggregate on the host with the interfaces connected to these ports. Both active-active and active-standby configurations for aggregate ports are allowed.

On each switch, create the LLA port-channel, ISL port-channel and add the member ports. The LLA port on each switch is the port connected to the host. The ISL ports are the ports connecting the two switches. The LLA port-channel and corresponding member ports must be the same on both switches. Then configure the LLA role as primary or secondary for both switches. Ensure that one switch is the primary and the other the secondary. Then enable LLA. If the `enable LLA` command is not executed, LLA will not be enabled after restart, even though the LLA role and ports are configured. Then restart the SEFOS on the switches. After the restart, LLA will be operable and the switches will assume their respective roles.

1. Create an LLA port-channel.

```
Switch-<n>(config)# interface port-channel 101
```

The LLA port-channels are reserved from 101 to 124. The mode changes to port-channel configuration.

```
Switch-<n>(config-if)# no shut  
Switch-<n>(config-if)# exit
```

2. Create the ISL port-channel.

```
Switch-<n>(config)# interface port-channel 100
```

The port-channel 100 is reserved for ISL. The mode changes to port-channel configuration.

```
Switch-<n>(config-if)# spanning-tree cost 10
```

The ISL port-channel must be in forwarding state on both the fabrics.

```
Switch-<n>(config-if)# no shutdown  
Switch-<n>(config-if)# exit
```

3. Add ports to the LLA port-channel.

```
Switch-<n>(config)# interface extreme-ethernet 0/15  
Switch-<n>(config-if)# channel-group 101 mode on  
Switch-<n>(config-if)# no shutdown  
Switch-<n>(config-if)# exit
```

4. Add ports to the ISL port-channel.

```
Switch-<n>(config)# interface extreme-ethernet 0/11
Switch-<n>(config-if)# channel-group 100 mode on
Switch-<n>(config-if)# no shutdown
Switch-<n>(config-if)# exit

Switch-<n>(config)# interface extreme-ethernet 0/12
Switch-<n>(config-if)# channel-group 100 mode on
Switch-<n>(config-if)# no shutdown
Switch-<n>(config-if)# exit
```

The ISL ports must be contiguous. The `interface range` command may also be used to add ports to the ISL port-channel.

5. Set the LLA role and add the ISL and LLA port-channels.

The role is configured as primary on one switch and as secondary on the other switch.

```
Switch-1(config)# lla
Switch-1(config-lla)# role primary

Switch-1(config)# lla
Switch-2(config-lla)# role secondary

Switch-<n>(config-lla)# isl port-channel 100
Switch-<n>(config-lla)# lla port-channel 101
Switch-<n>(config-lla)# exit
Switch-<n>(config)#
```

6. Enable LLA.

```
Switch-<n># configure terminal
Switch-<n>(config)# set lla enable

Switch-<n>(config)# set port-channel enable
```

The port-channel feature has to be enabled for LLA.

7. Verify the configuration.

```
Switch-1# show lla role

LLA Role: primary
Switch-1# show lla status
```

```
LLA Disabled with role not set at bootup
```

```
SEFOS. LLA configurations take affect only after saving the  
configuration and restarting
```

```
SEFOS# show lla ports
```

```
Isl Trunk Port-channel: port-channel 100
```

```
Ports:
```

```
interface extreme-ethernet 0/11
```

```
interface extreme-ethernet 0/12
```

```
Lla Port-channel: port-channel 101
```

```
Ports:
```

```
interface extreme-ethernet 0/15
```

8. Save the LLA configuration.

```
Switch-<n># write startup-config
```

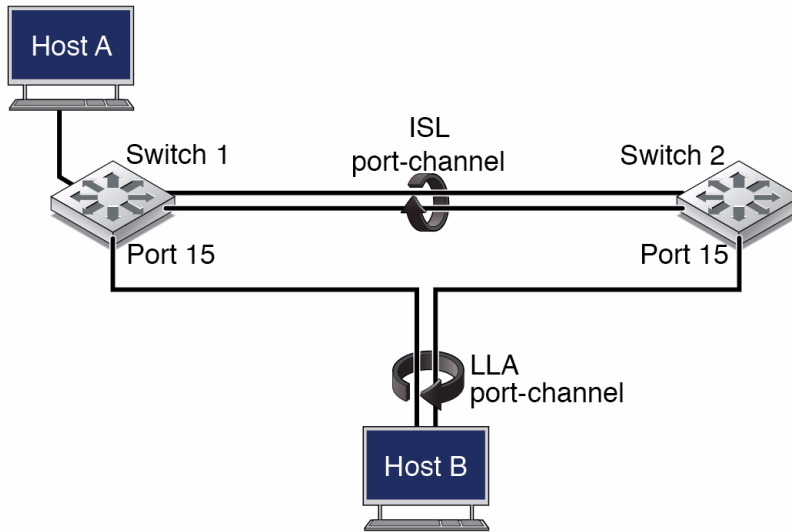
```
LLA configuration takes affect after restarting the SEFOS.
```

9. Verify the configuration by checking connectivity between a host connected to an uplink on one of the switches to another host connected to the LLA ports on both switches.

```
Switch-1# show lla status
```

```
LLA Enabled with role primary at bootup
```

- Ping from Host A on uplink port 10 on switch 1 to Host B on LLA ports 15. On Host B, the interfaces connected to the switches are bonded.



```
[Host-A ~]# ping host-b

PING host-b (192.168.1.10) 56(84) bytes of data.
64 bytes from host-b (192.168.1.10): icmp_seq=0 ttl=255 time=1.61
ms
64 bytes from host-b (192.168.1.10): icmp_seq=1 ttl=255 time=0.288
ms

--- host-b ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1002ms
rtt min/avg/max/mdev = 0.288/0.952/1.616/0.664 ms
```

- Shut down LLA port 15 on switch 1. Traffic resumes in very short time through port 15 on switch 2.

10. Show the running configuration for LLA.

```
Switch-1# show running-config LLA

Building configuration...
lla enabled
!
role primary
!
interface port-channel 100
isl trunk port-channel
!
interface extreme-ethernet 0/11
```

```
isl port
!  
interface extreme-ethernet 0/12  
isl port  
!  
!  
!  
interface port-channel 101  
lla port-channel  
!  
interface extreme-ethernet 0/15  
lla port  
!  
end
```