TECHNICAL NOTE
STORMSHIELD NETWORK SECURITY

LACP LINK AGGREGATION

Product concerned : SNS 1.0 and higher versions

Document version : 1.0

Reference : sns-en-LACP_link_aggregation_technical_note

Table of contents

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

Introduction

This document aims to guide Stormshield Network firewall administrators through the configuration and operation of the Link Aggregation Control Protocol (LACP), also known as IEEE802.3ad.

It will start by describing the protocol's operating principles and the hardware requirements for appliances. A typical LACP configuration between a Stormshield Network firewall and Cisco and HP switches will also be illustrated, along with troubleshooting information.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

Principles of link aggregation

Definition

Link aggregation refers to the grouping of several physically distinct interfaces into a single logical interface.

It serves several purposes:

- Ensuring fault tolerance in the event a link is down or issues arise on an interface,
- Increasing bandwidth between two interconnected appliances.

Link aggregation can be set up in various ways, generally in a parallel configuration and depending on the hardware used.

For example, between a router and a switch, certain mechanisms such as *Adaptive transmit load balancing* does not require any configuration on the switch, unlike LACP which requires both interconnected appliances to be configured.

Stormshield LACP - Prior requirements

LACP is supported from SNS v1.0 upwards. LACP in high availability is supported from SNS v2.0 upwards.

It can be set up on the following Stormshield Network firewalls:

NG1000, NG5000, SN510, SN710, SN910, SN2000, SN3000 and SN6000.

All aggregated physical ports must use the same settings:

- Speed,
- Duplex (half or full),
- 802.1q (unique VLAN ID or labeled multi-VLAN link).

Outbound traffic is sent over member interfaces of the aggregate link (maximum 8 interfaces).

LACP operates in two ways:

- Active mode: the appliance initiates the establishment of the aggregate.
- Passive mode: the appliance only responds to the LACP requests that it receives.

Stormshield Network firewalls only operate in active mode, so regardless of which mode (active or passive) the switches are in during LACP negotiation, the negotiation will be successful.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

Vendor implementations

EtherChannel by Cisco

Cisco switches are compatible with most standard link aggregators (e.g., LACP, PAgP, Bonding Round Robin, etc). The command prompt makes it possible to group several interfaces on a switch into an aggregate named EtherChannel (maximum 8 ports), then configure this aggregate's operating mode. All member ports of an EtherChannel aggregate must have the same speed.

Trunk by HP

HP switches use standardized LACP by default but enable the recognition of other types of link aggregates in a static configuration. This aggregate is called a "Trunk".

However, the term "Trunk" at Cisco refers to a link between two active elements that allows 802.1q tagged frames through for the identification of VLANs. So in order to avoid any ambiguity, it would be better to refer to link aggregation instead of Trunk or EtherChannel.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

Implementing the test workshop

Network topology used in this document

An aggregate of two interfaces has been defined on a Stormshield Network SN2000 firewall for each switch.

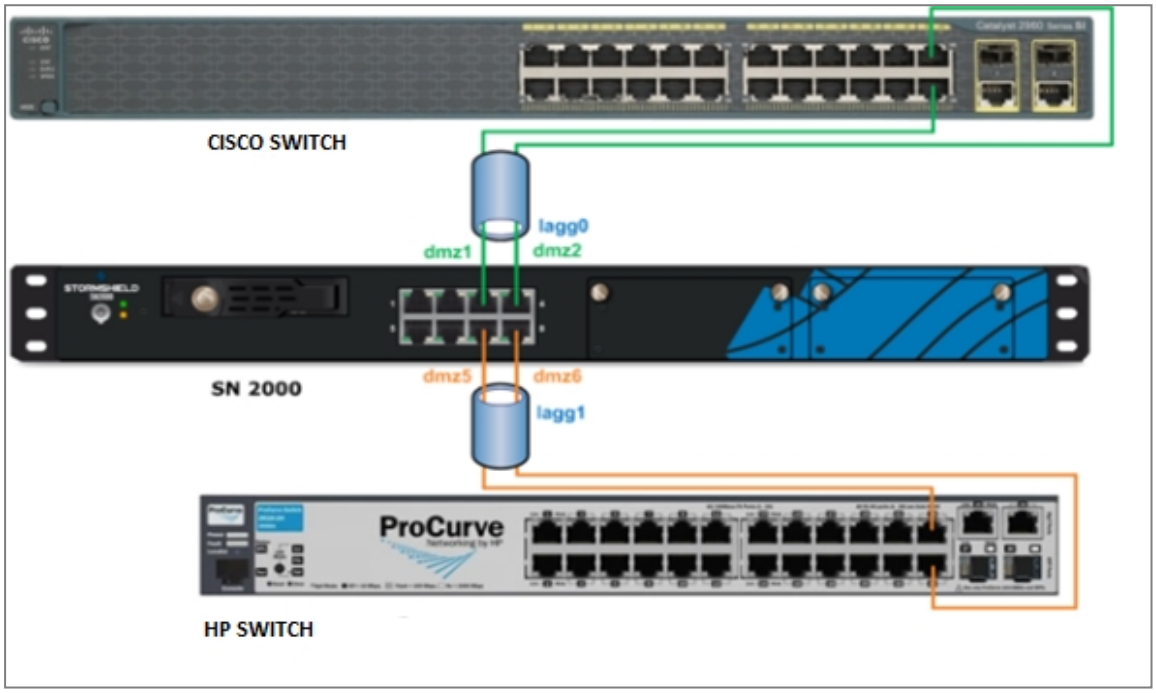

The following table provides the information relating to the interfaces of a Stormshield Network SN2000 firewall:

Name / number displayed (GUI)	Common name (network file)	Visual port number (appliance)
mgmt1 /(1)	Ethernet0	1 (rear)
mgmt2 /(2)	Ethernet1	2 (rear)
out /(3)	Ethernet2	1 (front)
in /(4)	Ethernet3	2 (front)
dmz1 /(5)	Ethernet4	3 (front)
dmz2 /(6)	Ethernet5	4 (front)
dmz3 /(7)	Ethernet6	5 (front)
dmz4 /(8)	Ethernet7	6 (front)
dmz5 /(9)	Ethernet8	7 (front)
dmz6 /(10)	Ethernet9	8 (front)

To find out the system name, use the `portinfo` command.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

Test handling

During the workshop, all the interfaces on the SN2000 appliance will belong to the factory bridge. Pings will be sent from a workstation on the network (Network_bridge) to the IP address of the bridge (firewall_bridge) via the firewall's aggregated interface.

A filter rule has been configured to allow ICMP Echo Request notifications.

Configuring the switch

Since the firewall is in active mode, the switch can remain in passive mode.

The detailed configuration in this example assumes the user's familiarity with the vendor's command line interface.

Only the link aggregation setup will be covered.

CISCO CLI	HP CLI	COMMENTS

`Switch(config)# interface range Gi 0/23-24`
 	
`Switch(config-if-range)# channel-group 1 mode passive`
 	
`Switch(config-if-range)# channel-protocol lacp`
	`ProCurve(config)# trunk 23-24 trk1 lacp`	Configuration of the aggregation in passive LACP mode on ports 23 and 24 in this example.

`Switch(config-if-range)# exit`
 	
`Switch(config)# interface Port-channel 1`
 	
`Switch(config-if)#switchport mode trunk`
 	
`Switch(config-if)#switchport trunk allowed vlan 2-10`
	`ProCurve(config)# vlan 2 tagged trk1`
	
	`ProCurve(config)# vlan 3 tagged trk1`
	…
	…
	…
	`ProCurve(config-vlan)# vlan 10 tagged trk1`
		Optional commands allowing 802.1q tagging on the logical switch-router link, if VLANs 2 to 10 are used in this example.

IMPORTANT

In order to avoid layer 2 issues (instability of the MAC address table, broadcast storm, etc.), configure the aggregate on the firewall and on the switch before interconnecting both appliances.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

NOTES

- After ensuring that the aggregate is running properly (by interrupting the link for example), you will need to back up the configuration of the switch.
- The naming and numbering of interfaces vary according to the chosen switch model. Refer to the vendor's user guide if necessary.
- The *Port-Security* feature found on Cisco and HP switches is not compatible with LACP, and should not be configured on any of the members in the aggregate.

Configuring the Stormshield Network firewall

Setting up a link aggregation on a Stormshield Network firewall will create an interface with the system name laggX, X being a number starting at zero.

The maximum number of aggregates possible on a Stormshield Network firewall is equivalent to N/2 (where N is the number of Ethernet ports).

Configuring the firewall via the administration interface

1. In the interface configuration window, select the first interface to be aggregated, then click on **Add > Convert to link aggregation (LACP)**.

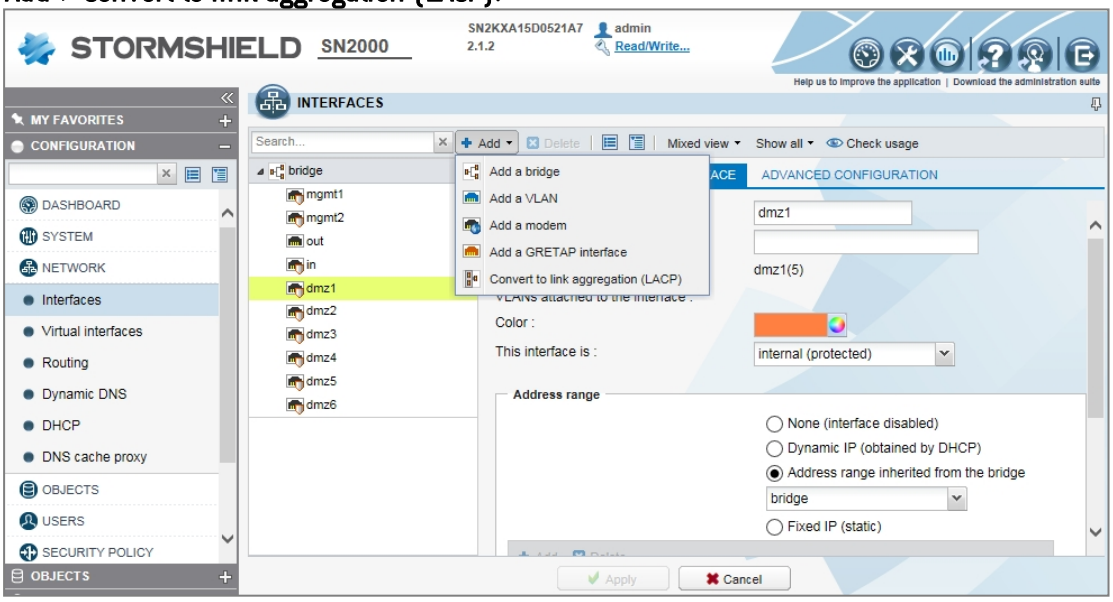

2. In the dialog box that appears, click on **Convert** to confirm.
3. In the *Configuration of the interface* tab, rename the interface in the **Name** field.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

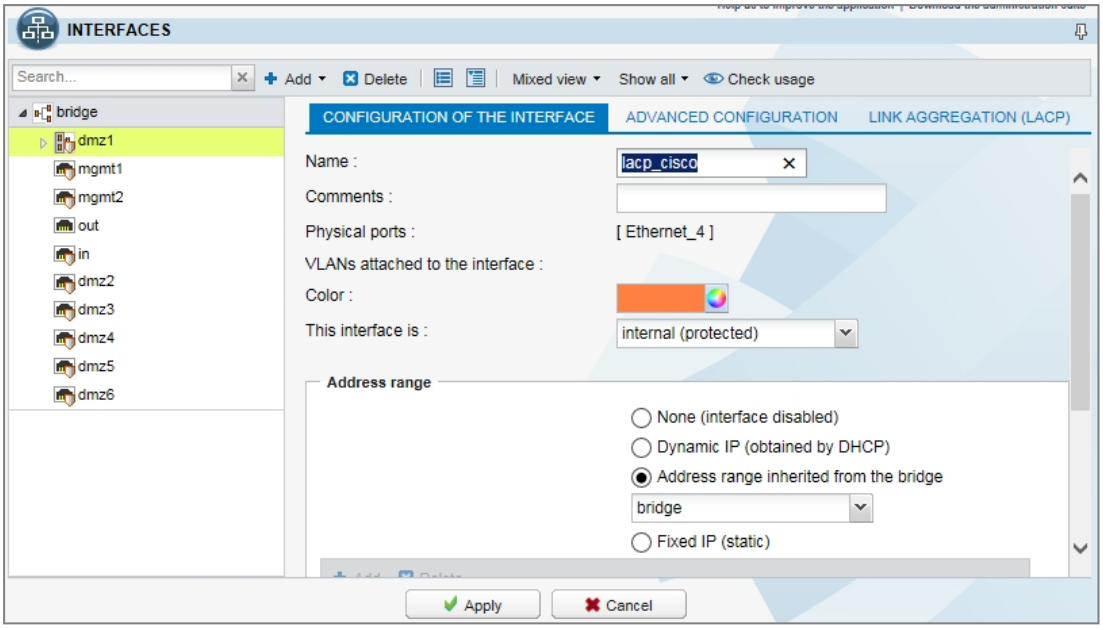

4. In the *Link aggregation (LACP)* tab, add an interface to the aggregate.

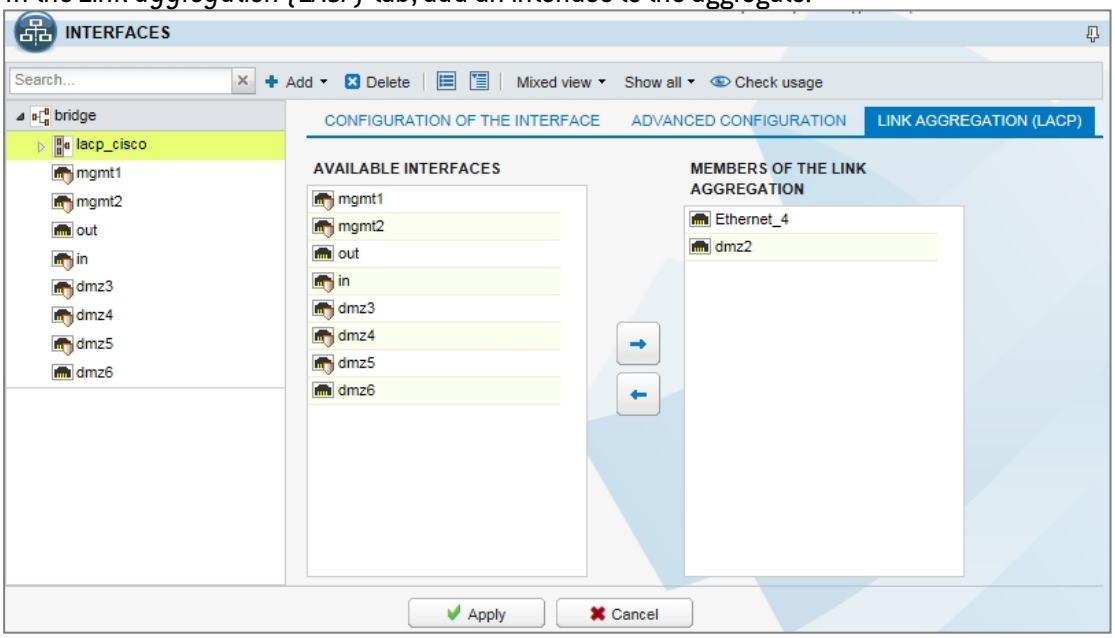

5. Check that the configuration has been applied.

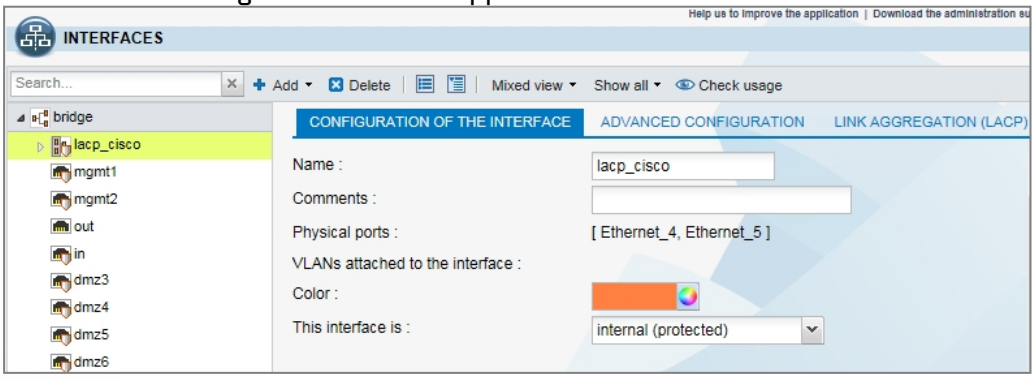

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

6. Repeat the same operation for aggregated interfaces on the HP switch side.

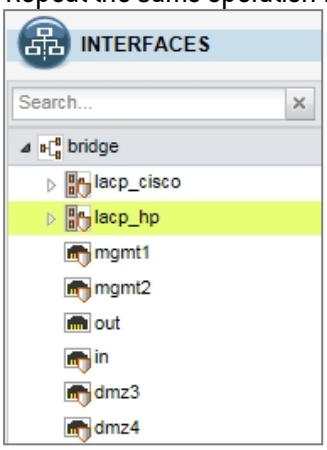

7. If necessary, check the configuration of a member interface in an aggregate.

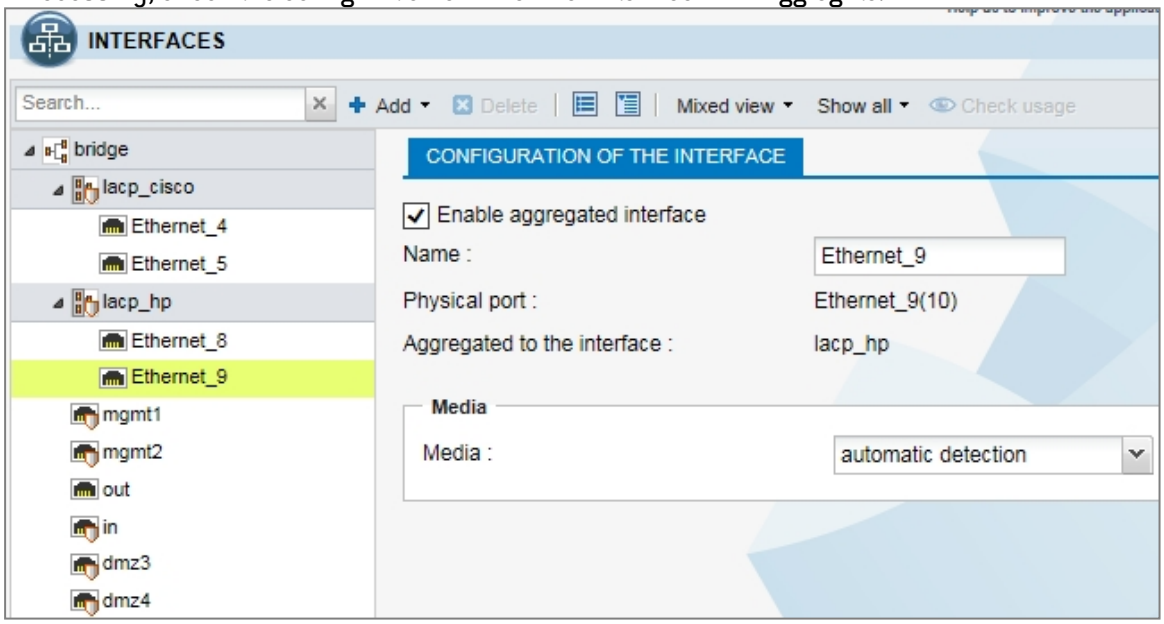

Verifying the configuration

SN2000 CLI	Comments

`SN2KXA15D0521A7>portinfo`
`port name NS-BSD state addressIPv4 addressIPv6`
`1 mgmt1 igb8 no-link 10.10.30.193/27`
`2 mgmt2 igb9 no-link 10.10.30.193/27`
`3 out igb4 no-link 10.10.30.193/27`
`4 in igb5 up 10.10.30.193/27`
`5 lacp_cisco lagg0/igb6 no-link 10.10.30.193/27`
`6 lacp_cisco lagg0/igb7 no-link 10.10.30.193/27`
`7 dmz3 igb0 no-link 10.10.30.193/27`
`8 dmz4 igb1 no-link 10.10.30.193/27`
`9 lacp_hp lagg1/igb2 no-link 10.10.30.193/27`
`10 lacp_hp lagg1/igb3 no-link 10.10.30.193/27`
	With the `portinfo` command, ports are referred to by the numbering shown in the interface.
	
	Take note of the interfaces' NS-BSD names, which will be used in logs.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

SN2000 CLI	Comments

SN2KXA15D0521A7>ifinfo
interface list:
bridge
  10.10.30.193/27
  out (igb4)
  dmz3 (igb0)
  dmz4 (igb1)
  lacp_cisco (protected,lagg0) { igb7, igb6, }
  lacp_hp (protected,lagg1) { igb3, igb2, }
  mgmt1 (protected,igb8)
  mgmt2 (protected,igb9)
  in (protected,igb5)
sslvpn (protected,tun0)
loopback5 (lo5)
loopback4 (lo4)
loopback3 (lo3)
loopback2 (lo2)
loopback1 (lo1)
ipsec (enc0)
	The NS-BSD names of aggregated interfaces will appear in the description of aggregates.

Now that the configuration of aggregates is complete both on the Stormshield Network firewall and on the switches, Ethernet cables can be connected.

Diagnosing the switch

CISCO CLI	HP CLI
Switch#show interface port-channel 1 Port-channel1 is up, line protocol is up (connected) Hardware is EtherChannel, address is 001d.4608.e217 (bia 001d.4608.e217) MTU 1500 bytes, BW 2000000 Kbit, DLY 10 usec, Certain lines have been omitted here but appear on the screen.	No HP equivalent

Comments: Take note of the bandwidth on the Cisco side: BW = 2000000 Kbit (2x1 Gbps).

CISCO CLI	HP CLI

Switch#show etherchannel 1 detail
Group state = L2
Ports: 2 Maxports = 8
Port-channels: 1 Max Port-channels = 8
Protocol: LACP
Minimum Links: 0
 	
Certain lines have been omitted here but appear on the screen.
 	
Ports in the Port-channel:
Index Load Port EC state No of bits
-----+----+-----+--------+-------
  0 00 Gi0/23Passive 0
  0 00 Gi0/24Passive 0
 	
Certain lines have been omitted here but appear on the screen.
	ProCurve(config)#show lacp
	
	  LACP
	PORT LACP TRUNK PORT LACP
	NUMB ENABLED GROUP STATUS PARTNER
	1 Passive 1 Down No
	
	2 Passive 2 Down No
	
	Certain lines have been omitted here but appear on the screen.
	
	23 Active Trk1 Up Yes
	24 Active Trk1 Up Yes

Comments: LACP is associated with the chosen aggregation mode. Warning: The term `Passive` shown in the HP CLI refers to the interface and does not concern the LACP configuration mode.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

CISCO CLI	HP CLI
Switch#debug etherchannel PAgP/LACP Shim/FEC debugging is on **00:14:44: FEC: Un-Bndl msg NOT send to PM for port Gi0/23 from Po1** 00:14:44: FEC: delete port (Gi0/23) from agport (Po1) 00:14:44: FEC: lacp_switch_remove_port_from_ associated_list_internal: Gi0/23 deleted from the associated list for Po1 00:14:45: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/23, changed state to down Certain lines have been omitted here but appear on the screen. 00:59:50: %LINK-3-UPDOWN:**Interface GigabitEthernet0/23, changed state to up** Certain lines have been omitted here but appear on the screen. 00:59:51: FEC: add port (Gi0/23) to agport (Po1) 00:59:52: FEC: lacp_fec_bundle_internal: Determine if msg to PM to bundle port Gi0/23 with Po1 is needed 00:59:52: FEC: pagp_switch_want_to_bundle: Bndl msg to PM for port Gi0/23 to Agport Po1 00:59:52: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/23, changed state to up	ProCurve(config)#show log Keys: W=Warning I=Information M=Major D=Debug --- Event Log listing: Events Since Boot --- I 13/10/15 00:00:57 ports: trunk Trk1 is now active I 13/10/15 00:00:57 ports: port 23 in Trk1 is now on-line I 13/10/15 00:00:57 ports: port 24 in Trk1 is now on-line Certain lines have been omitted here but appear on the screen. **I 13/10/15 09:52:58 ports: port 23 in Trk1 is now off-line** **I 13/10/15 09:52:58 ports: port 25 is Blocked by LACP** **I 02/10/13 09:52:58 ports: port 23 in Trk5 is now on-line**

Comments: Pings are sent to the SN firewall from a workstation connected to the switch.

The cable connected to the first port in the aggregate has been unplugged.

Cisco's debug mode or HP's log mode shows changes that have been made.

When the cable is plugged back in, the aggregate will switch back to its usual status.

Don't forget to exit debug mode on the Cisco switch.

While no pings are ever lost when they pass through the HP switch, they will be on the Cisco switch every time its status changes (when one of the cables is unplugged or plugged back in).

Diagnosing the Stormshield Network firewall

SN2000 commands	Comments
	Pings are sent to the SN firewall from workstations connected to Cisco and HP switches.
SN2KXA15D0521A7>grep -i "alarmid=3| alarmid=8" /log/l_alarm	Alarm IDs relating to LACP begin with a 3 or 8.
id=firewall time="2015-11-09 14:47:49" fw="SN2KXA15D0521A7" tz=+0000 startime= "2015-11-09 14:47:49" pri=4 msg="Lien agrégé activé: igb6" class=system alarmid=84	Connection of link 1 to Cisco: dmz1 alias igb6 (as seen in the results of the portinfo command).
id=firewall time="2015-11-09 14:47:49" fw="SN2KXA15D0521A7" tz=+0000 startime= "2015-11-09 14:47:49" pri=4 msg="Interface activée: lagg0" class=system alarmid=37	The aggregation to Cisco only works on link 1; pings will pass through.
id=firewall time="2015-11-09 3:08:05 PM" fw="SN2KXA15D0521A7" tz=+0000 startime= "2015-11-09 3:08:05 PM" pri=4 msg="Lien agrégé activé: igb7" class=system alarmid=84	Connection of link 2 to Cisco: dmz2 alias igb7.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

SN2000 commands	Comments
id=firewall time="2015-11-09 15:08:18" fw="SN2KXA15D0521A7" tz=+0000 startime= "2015-11-09 15:08:18" pri=4 msg="Lien agrégé désactivé: igb6" class=system alarmid=85	Disconnection of link 1; pings continue to pass through.
id=firewall time="2015-11-09 3:08:35 PM" fw="SN2KXA15D0521A7" tz=+0000 startime= "2015-11-09 3:08:34 PM" pri=4 msg="Lien agrégé activé: igb6" class=system alarmid=84	Reconnection of link 1.
id=firewall time="2015-11-09 15:10:12" fw="SN2KXA15D0521A7" tz=+0000 startime= "2015-11-09 15:10:11" pri=4 msg="Interface désactivée: lagg0" class=system alarmid=38	Disconnection of both links, the interface lagg0 is disabled (no more pings going through the Cisco aggregate).
id=firewall time="2015-11-09 3:10:17 PM" fw="SN2KXA15D0521A7" tz=+0000 startime= "2015-11-09 3:10:17 PM" pri=4 msg="Lien agrégé activé: igb2" class=system alarmid=84	Connection of link 3 to HP: dmz5 alias igb2.
id=firewall time="2015-11-09 3:10:17 PM" fw="SN2KXA15D0521A7" tz=+0000 startime= "2015-11-09 3:10:17 PM" pri=4 msg="Interface activée: lagg1" class=system alarmid=37	The aggregation to HP only works on link 3.
id=firewall time="2015-11-09 15:11:04" fw= "SN2KXA15D0521A7" tz=+0000 startime= "2015-11-09 15:11:04" pri=4 msg="Lien agrégé activé: igb3" class=system alarmid=84	Connection of link 4 to HP: dmz6 alias igb3.

In brief, when a link in the aggregate is connected, a system alarm id=84 will be raised, while a system alarm id=37 indicates that the aggregate has become operational.

If all links fail in an aggregate, a system alarm id=38 will indicate that the aggregate itself has been disabled, whereas a single disabled link will raise the system alarm id=85.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

Going a little further...

Enabling verbose mode in LACP

1. Log on to the firewall in SSH and run the following commands:

```
cp /etc/syslog.conf /etc/syslog.conf.bak
touch /log/lacp.log
```

2. Edit the file */etc/syslog.conf* and modify the line `kern.*` by specifying the path of the file created earlier:

```
kern.*                          /log/lacp.log
```

3. Change the log level in order to enable the LACP verbose mode:

```
sysctl net.link.lagg.lacp.debug=3
```

4. Run the logging process once more:

```
killall -HUP syslogd || syslogd
```

LACP logs in verbose mode can be accessed in `/log/lacp.log`.

To disable verbose mode, log on to the firewall in SSH and run the following commands:

```
sysctl net.link.lagg.lacp.debug=0
mv /etc/syslog.conf.bak /etc/syslog.conf
```

Link aggregation limitations

All physical ports in the same aggregate must appear on a single switch, or even on a single ASIC on this switch, as is the case on an HP Procurve. An ASIC (Application Specific Integrated Circuit) represents a port grouping here.

However, for stackable switches, traffic may be balanced on separate switches.

For example, on the Cisco 3750 switch range, the technology is called Cross-Stack EtherChannel and setting it up requires two switches in a VSS (Virtual Switching System) configuration, allowing both physical switches to run as a logical switch.

Another example is the proprietary protocol DMLT (Distributed Multi-Link Trunking) developed by Avaya.

A highly resilient infrastructure can therefore be implemented, but with significant capital investment on layer 2 hardware (switches), and in terms of administration time. Examples of fault-tolerant infrastructures will be given further on in this document.

SNS HA clusters and LACP

This section does not explain how to configure a HA cluster.

Calculating the HA quality factor with LACP

You may recall that the quality factor is a mathematical formula that depends on the status (and sometimes the weight) of active interfaces, calculated on both nodes of the cluster in order to nominate the active node.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

When the presence of an aggregate is detected, the calculation will be based on the status of the aggregated interface instead of the member interfaces of the aggregate so that the loss of a single interface in the aggregate would not change the quality factor.

Example

The SN2000 appliance used earlier has 10 active Ethernet interfaces of equal weight by default.

All of the interfaces have a working link.

An aggregate *lagg0* contains two aggregated interfaces: 8 single interfaces and an aggregate, i.e., 9 interfaces.

The default weight is 100.

With this configuration, the quality factor will be calculated as follows:

(9x100 / 9x100)x100 = 100%

If either of the aggregated interfaces loses its connection, the quality factor will remain the same as it remains in the aggregate of an active interface.

If the aggregate loses its connection, the quality factor will change as follows:

(8x100 / 9x100)x100 = 89%.

NOTE

Even in SSH, this behavior cannot be modified.

Use case 1: one Stormshield Network cluster, one switch

Recommended topology

The HA cluster is equipped with two control links on dedicated interfaces.

For switches connected to the cluster, create an aggregate for each node of the cluster to be linked.

Configuring the cluster with LACP

In the event of a swap (modification of the cluster's active node), the LACP renegotiation initiated by the new active node may take about 30 seconds, preventing any traffic activity on the aggregate.

The best way to work around this renegotiation is to configure the LACP negotiation on both nodes of the cluster according to the following procedure, corresponding to your SNS version:

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

SNS from v3 upwards

1. In the active firewall's administration interface, select **System > High availability**.
2. In the **Swap configuration** area, select the checkbox **Enable link aggregation when the firewall is passive**.
3. In the upper banner of the interface, click on the Cluster icon to synchronize both firewalls.

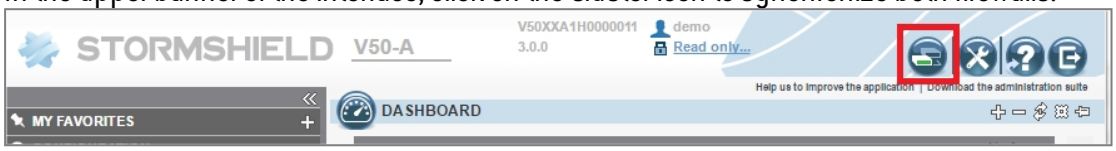

The LACP aggregate will then be configured on both nodes.

SNS v2 (not available in v1)

1. Log on to the active node in SSH and run the following command:
```
setconf ~/ConfigFiles/HA/highavailability Global LACPWhenPassive 1
```
2. Synchronize the configuration on the passive node by running the following commands:
```
enha -f -v
hasync -v
```

The LACP aggregate will then be configured on both nodes.

Use case 2: one Stormshield Network cluster, two stackable switches

Recommended topology

Logically, the configuration diagram would be the same as in the previous case, since both switches are considered a single logical switch. However, the switch is no longer a SPOF (Single Point Of Failure).

The following physical configuration diagram will therefore enable full fault tolerance:

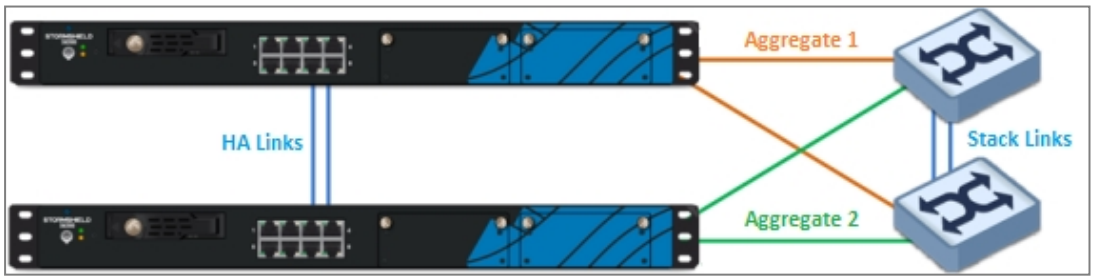

Two Stormshield Network firewalls form a high availability cluster with an active node and a passive node.

Each node in the cluster is linked to two separate switches through link aggregation.

Both switches are stackable. Depending on brands and models, there may be an active switch and a passive switch, or both switches may be in active mode. Switches make it possible to spread out aggregates as two aggregated interfaces can be spread out over two physical switches.

Configuring the cluster and switches with LACP

Configure the cluster in the same way it was defined in the previous use case. For more information, please refer to the section Configuring the cluster with LACP.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

In the network diagram below (in which workstations are shown), the issue of network looping may arise.

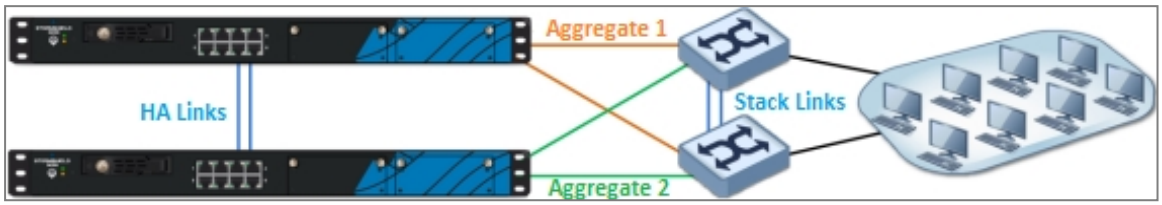

Loops appear in a network when several different level 2 links (data link layer in the OSI model) exist between two active elements.

Stackable switches have to be equipped with a loop-free system in order to prevent stability issues in the MAC address table, or operate in active-passive mode.

However, no changes need to be made in the SNS cluster: whether an Ethernet frame containing the source MAC address of a workstation passes through switch 1 or switch 2, it will be presented to the cluster on the aggregated interface of the active node.

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016

STORMSHIELD

documentation@stormshield.eu

sns-en-LACP_link_aggregation_technical_note - Copyright © Stormshield 2016