



Nortel Networks L2/3 GbESM Layer 2 High Availability Enhancements – L2/3 Switch Software Version 1.1

Planning / Implementation (withdrawn product)

Layer 2 High Availability Enhancements – L2/3 switch software version 1.1

Introduction

The “trunk failover” feature in this software release includes additional functionality which makes more versatile and granular High Availability configurations possible. These enhancements are in the following areas:

- Support for trunk failover with LACP trunks as well as static trunks
- Increased granularity for trunk failover – specific VLAN(s) can be failed over instead of all or nothing failover
- It is now possible to specify a minimum number of active links below which failover will occur instead of having failover occur only when all links enabled for failover support are down.
- Up and down for purposes of trunk failover now relate to the STP forwarding (FWD) state of the port rather than its link state.

All of this is configured via a new `/cfg/l2/failovr` menu, which is described in the sections that follow.

Applications

In order to decide whether and how to use these enhancements, consider the following questions:

1. Do you need to selectively control internal port failover:
 - a. to fail over some of the internal ports, but not others, to the other switch?
 - b. to fail over some or all internal ports in the event that an identified set of external ports fail?
 - c. to trigger failover when the number of active external links drops below a configured limit?
2. Which external ports are associated with which internal (server blade) ports:
 - a. What is the minimum number of external links that can support the application successfully (or the number below which failover should be triggered)?

These enhancements can be useful in circumstances such as the following:

1. Customers who have split the ownership of a single BladeCenter chassis can use the ability to segment failover based on which trunk failed and cause only the associated blades to fail over to the backup switch. Split ownership can mean merely that two distinct applications are running in the chassis or it can mean having two different organizational units own different ports. A configuration useful for this case is shown in Example 3.
2. Customers who have applications which are sensitive to the amount of bandwidth available can use the ability to cause a failover when some uplink ports have failed. A configuration useful for this case is shown in Example 2.
3. Any customers using LACP will benefit from the ability to use trunk failover on LACP trunks (also shown in Example 3).
4. All customers will benefit from the software using STP “FWD” state to identify those ports which are up; this reduces issues resulting from link flapping.

What has not changed

These software changes are for the trunk failover feature and not for VRRP Hot Standby. Use of NIC teaming with Hot Standby is still not supported.

In addition, it is still necessary to ensure that the NIC teaming drivers on the server blades are configured consistently with the way this feature is configured on the switch.

Glossary of terms

Link Aggregation Control Protocol (LACP)	A standards based protocol for dynamic configuration, setup, and tear-down of multi-link trunks between two network devices.
Monitored ports	A set of ports configured as part of a trigger whose STP state is monitored to determine when failover should occur.
Trigger	A configured item which identifies either one or more statically defined trunks (/cfg/l2/trunk) or one LACP key (/cfg/l2/lacp). Internal ports that are associated with the same VLANs as the monitored ports in the trigger are brought down when the trigger fires.
VLAN association	All of the internal ports which are members of one or more of the same VLANs as the monitored ports for a particular trigger are associated with that trigger. These ports are brought down when the trigger fires.

Configuration menus: /cfg/l2/failovr and submenus

The existing failover **ena** command is replaced by the following menus. Existing configurations are migrated seamlessly to the new syntax.

```
-----
[Failover Menu]
trigger - Trigger Menu
vlan - Globally turn VLAN Monitor ON/OFF
on - Globally turn Failover ON
off - Globally turn Failover OFF
cur - Display current Failover configuration
```

The **vlan** option allows the user to perform trunk failover with or without VLAN distinction. If the **vlan** option is **off**, only one Trigger can be enabled. No VLAN distinction is made, and all internal ports will be brought down when the trigger condition is satisfied. If the **vlan** option is **on**, multiple triggers can be enabled, where each trigger operates only on the set of VLANs associated with the trunks being monitored.

Configurations that are generated with software release 1.0.x.x are migrated automatically to the new software with the **vlan** option set to **off** (Example 1).

Configure failover trigger: /cfg/l2/failovr/trigger

Failover parameters that are related to static trunks and LACP trunks can be specified in the menu /c/l2/failovr/trigger <1-8>. Up to eight triggers can be defined and enabled individually.

Note: Failover configurations that generated using Release 1.0.x.x are migrated as follows:

- all failover enabled trunks are added to the monitor for Trigger 1
- limit is set to 0 (for example, all ports in Trigger 1 must go down before failover is triggered)

```
>> Failover# trigger
Enter Trigger ID: (1-8)
```

```
-----
[Trigger 1 Menu]
amon - Auto Monitor Menu
limit - Limit operational links
ena - Enable Trigger
dis - Disable Trigger
cur - Display current Trigger configuration
```

Configure trigger limit: /cfg/l2/failovr/trigger <#>/limit

The limit option specifies the minimum number of operational links within a trigger that would initiate a failover event. For example, a limit of 2 means a failover event occurs if the number of operational links in the trigger is 2 or less (for example, three links must be active to avoid failover). To initiate failover only when all links in the trigger fail (i.e. no operational links are available), set the limit option to 0.

Configure auto monitor: /cfg/l2/failovr/trigger <#>/amon

The Auto Monitor Menu defines what to monitor: static trunks or an LACP trunk that is referenced by the LACP admin key. If the vlan option is on, the internal control ports are selected automatically based on the VLANs that are associated with the ports that are monitored. Otherwise, all the internal ports are selected regardless of the VLAN group.

The CLI configuration validation code enforces the following restrictions:

- Multiple triggers are NOT allowed if the vlan option is OFF.
- Multiple triggers must NOT operate on the same internal port.
- Each trigger may monitor multiple static trunks or a single LACP key, but not both.
- All external ports in a static trunk must belong to the same VLAN membership and share the same PVID.
- All LACP ports with the same admin key must belong to the same VLAN membership and share the same PVID.
- Multiple triggers must NOT operate on the same VLAN selections.
- If multiple static trunks are added to the same trigger, all trunks within the trigger must belong to the same VLAN membership and share the same PVID.
- If a trigger operates on multiple VLANs which are not all in the same STG group, the failover monitor will check the port STP state on the default PVID only.

```
>> Trigger 1# amon
```

```
-----
[Auto Monitor Menu]
addtrnk - Add trunk to Auto Monitor
remtrnk - Remove trunk from Auto Monitor
addkey - Add LACP port adminkey to Auto Monitor
remkey - Remove LACP port adminkey from Auto Monitor
cur - Display current Auto Monitor configuration
```

Configuration examples

Example 1: Provide the equivalent function to that available in release 1.0.x.x

Features used: no new features.

Assumptions:

- Monitor trunks 1 and 2 and trigger failover when they both fail. This is equivalent to:

```

/cfg/l2/trunk 1/failovr ena
/cfg/l2/trunk 2/failovr ena

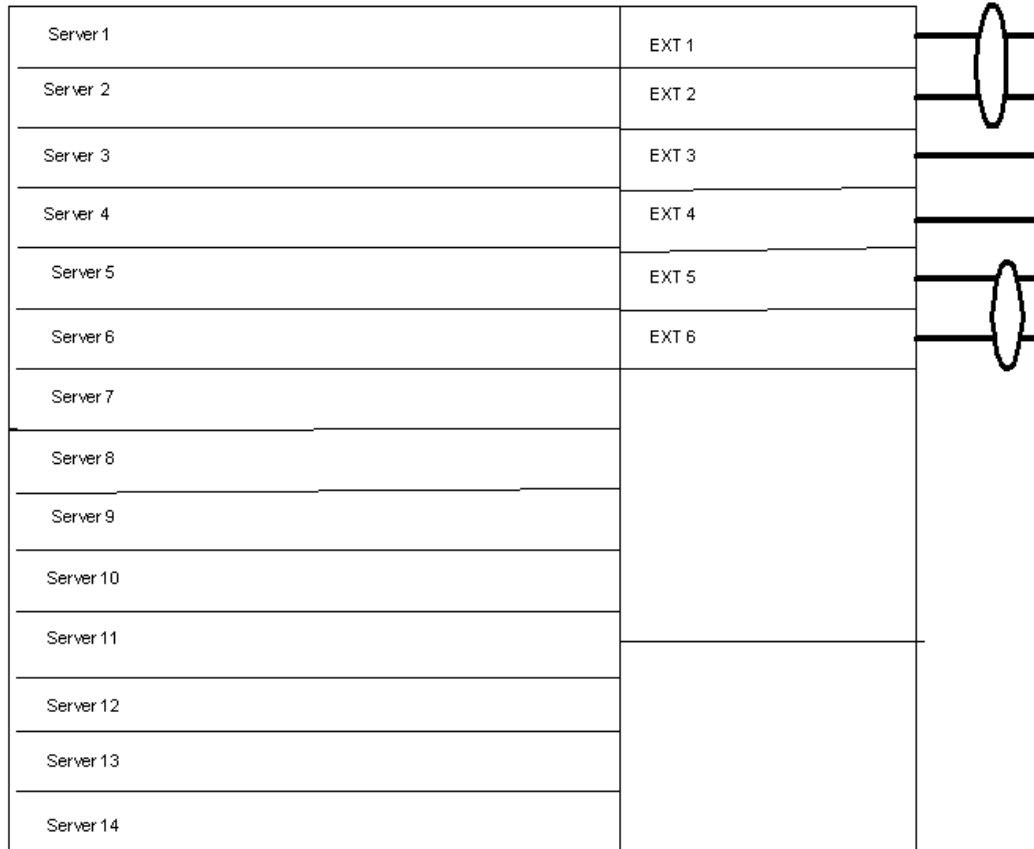
```

Configuration text:

```

/cfg/l2/failovr/on
/cfg/l2/trigger 1
ena
/cfg/l2/trigger 1/amon
addtrnk 1
addtrnk 2
_

```



Example 1: All 14 Server ports will be brought down when all ports in Trunk 1 and Trunk 2 are down

Example 2: Failover triggered by partial uplink failure

Features used:

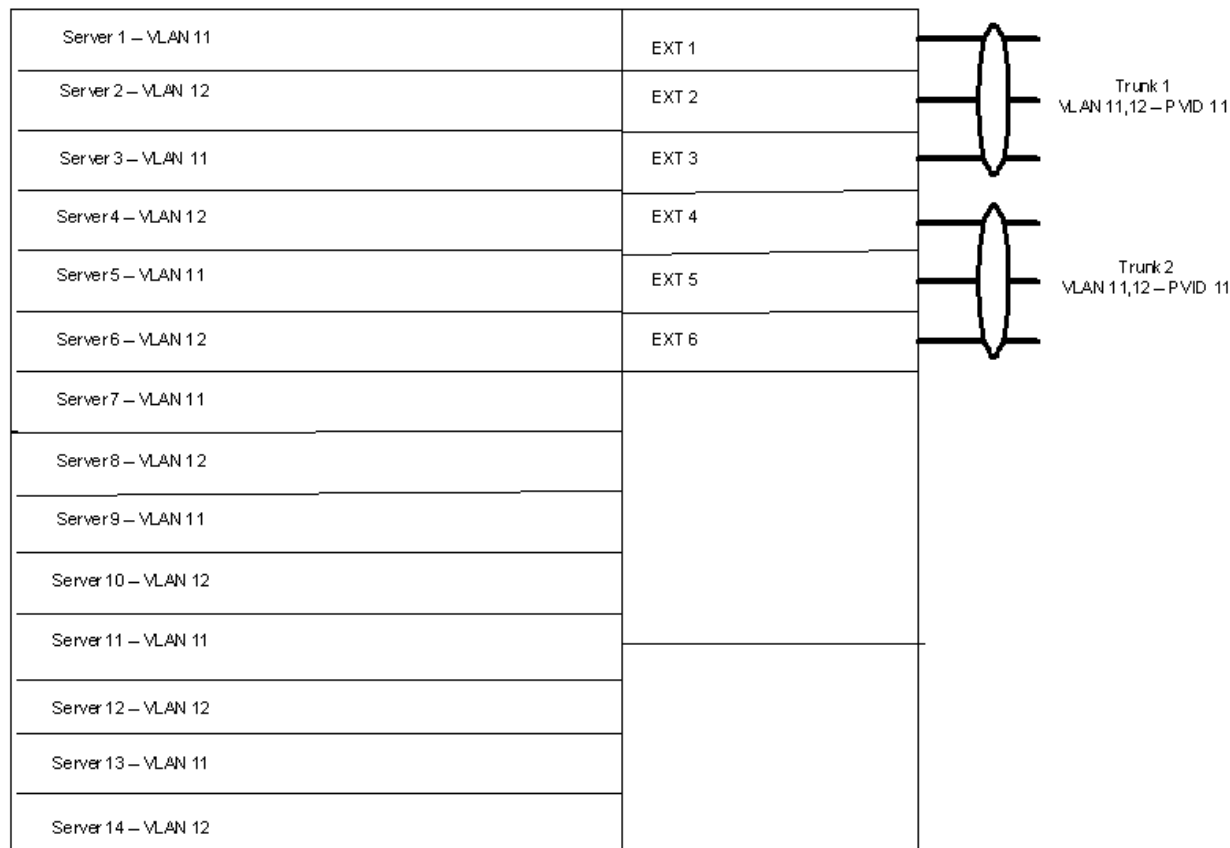
- limit command to trigger failover when some links are still active.

Assumptions:

- all internal ports are on VLAN 11 or 12
- all external ports are on VLANs 11 AND 12 with PVID 11
- trunk 1 is ports EXT1-3; trunk 2 is ports EXT4-6
- Failover is to be triggered when two or fewer ports are active (i.e. when any two of the four configured ports have failed regardless of their trunk membership)

Configuration text:

```
/cfg/l2/failovr/on
/cfg/l2/failovr/trigger 1
ena
limit 2
/cfg/l2/failovr/trigger 1/amon
addtrnk 1
addtrnk 2
_
```



Example 2: All 14 Server ports will be brought down when any two or more ports from Trunks 1 and 2 are down

Example 3: “Split Ownership” of a BladeCenter chassis

Assumptions

- Internal ports 1-7 and external ports 1-3 owned by one application or organizational unit; Internal ports 8-14 and external ports 4-6 owned by another.
- Internal 1-7 spread across vlans 11, 12; EXT 1-3 are members of trunk 1; they are all tagged carrying VLAN 11 and 12 with PVID 11.
- Internal 8-14 spread across vlans 21,22; EXT 4-6 use LACP key 20 and are tagged carrying VLAN 21 and 22 with PVID 21.

Features used:

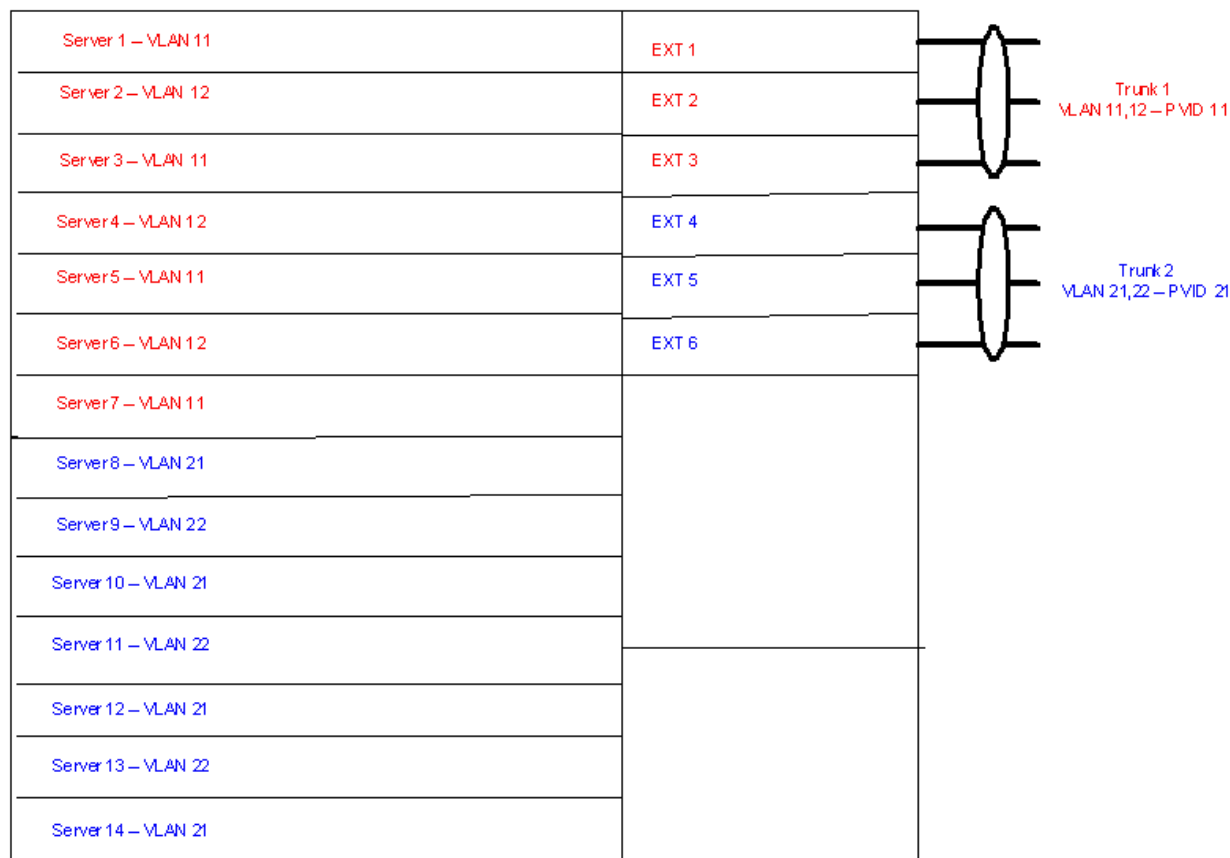
- Multiple triggers and the **vlan on** command necessary to enable their use.
- Monitoring an LACP trunk using the **addkey** command in a trigger.

Configuration text:

```
/cfg/l2/failovr/on
/cfg/l2/failovr/vlan on <- note that this is required because there are two triggers
/cfg/l2/failovr/trigger 1
ena
/cfg/l2/failovr/trigger 1/amon
addtrnk 1
/cfg/l2/failovr/trigger 2
ena
/cfg/l2/failovr/trigger 2/amon
addkey 20 <- this trigger monitors an LACP trunk with adminkey 20
```

Notes:

- The INT ports, which will be disabled in the event of failover for each trigger, are those in the same VLAN(s) as the members of the trunk. This means that ports INT1-7 will be disabled by trigger 1 whether they are members of VLAN 11 or VLAN 12. Similarly, ports INT8-14 will be disabled by trigger 2. The choice of internal ports is automatic based upon VLAN membership and cannot be overridden.
- Either one LACP key or one or more static trunks can be used in a trigger without using the **vlan on** command.
- If any internal port was configured to be a member of VLAN 11 and 21 (for example), which would cause it to be disabled by triggers 1 and 2, then the configuration would be invalid and would not apply successfully.
- Any internal ports which are not members of any VLAN which is configured on an external port will never be disabled by a trigger. For example, if INT1 is a member of VLAN 16 and relies on L3 forwarding to leave the chassis on EXT1-3, it will never be disabled. In such cases, L3 failover techniques such as VRRP should be used.



Example 3: Internal ports 1-7 will be brought down when
EXT1-3 are all down. Internal ports 8-14 will be brought down
when EXT4-6 are all down.

Example 4: Invalid configuration with overlapping VLANs

This example shows configuration options which would be rejected with diagnostic messages.

Assumptions:

- INT 1,3,5,7 on VLAN 11
- INT 2,4,6 on VLAN 12
- INT 9,11,13 on VLAN 21
- INT 8,10,12,14 on VLAN 22
- EXT 1,2 are members of trunk 1, carrying VLAN 11,12 PVID 11
- EXT 3,4 are members of trunk 2, carrying VLAN 21,22 PVID 21
- EXT 5,6 are members of trunk 3, carrying VLAN 11,12,21,22 PVID 11

Configuration text:

```
/cfg/l2/failovr/on
/cfg/l2/failovr/trigger 1
ena
/cfg/l2/failovr/trigger 1/amon
addtrnk 1
addtrnk 3 < -- this generates an error because EXT5 and EXT6 have different VLAN membership than EXT1 and 2
/cfg/l2/failovr/trigger 2
ena
/cfg/l2/failovr/trigger 2/amon
addtrnk 2
addtrnk 3 < -- this generates an additional error
because trunk 3 is configured for
two different triggers
/cfg/l2/failovr/trigger 3
ena
/cfg/l2/failovr/trigger 3/amon
addtrnk 3 < -- this will generate an error EVEN IF
trunk 3 is removed from the other two
triggers because of VLAN overlap with
both trigger 1 and 2
```

Notes:

The problems that this configuration illustrates could be corrected by making one of the following changes:

- Turning the VLAN sensitivity OFF and disabling all ports when an insufficient number of links are active to use only one trigger with trunks 1, 2, and 3.
- Reconfiguring ports 5 and 6 to carry either VLANs 11 and 12, or 21 and 22. Trunk 3 would then be only in the one trigger that is associated with the VLANs it carried. Trigger 3 would be removed.
- Allocating port 5 to trunk 1 and port 6 to trunk 2 and configuring them with the same VLAN membership as the other ports in those trunks. Trigger 3 and trunk 3 would be removed in this case.

Related product families

Product families related to this document are the following:

- [1 Gb Embedded Connectivity](#)
- [Blade Networking Modules](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2025. All rights reserved.

This document, TIPS0597, was created or updated on March 8, 2006.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/TIPS0597>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/TIPS0597>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

BladeCenter®

The following terms are trademarks of other companies:

Redbooks® is a trademark of IBM in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.