



MPLS VPN — OSPF and Sham-Link Support

Feature History

Release	Modification
12.2(8)T	This feature was introduced.
12.0(21)ST	This feature was integrated into Cisco IOS Release 12.0(21)ST, and support for Cisco 12000 series Internet Routers was added.
12.0(22)S	This feature was integrated into Cisco IOS Release 12.0(22)S, and support for Cisco 10000 series Internet Routers was added.
12.2(14)S	This feature was integrated into Cisco IOS Release 12.2(14)S.

This document describes how to configure and use a sham-link to connect Virtual Private Network (VPN) client sites that run the Open Shortest Path First (OSPF) protocol and share backdoor OSPF links in a Multiprotocol Label Switching (MPLS) VPN configuration.

This document includes the following sections:

- [Feature Overview, page 2](#)
- [Supported Platforms, page 8](#)
- [Supported Standards, MIBs, and RFCs, page 9](#)
- [Prerequisites, page 9](#)
- [Configuration Tasks, page 9](#)
- [Configuration Examples, page 11](#)
- [Command Reference, page 12](#)
- [Glossary, page 13](#)



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2007 Cisco Systems, Inc. All rights reserved.


Feature Overview

Using OSPF in PE-CE Router Connections

In an MPLS VPN configuration, the OSPF protocol is one way you can connect customer edge (CE) routers to service provider edge (PE) routers in the VPN backbone. OSPF is often used by customers that run OSPF as their intrasite routing protocol, subscribe to a VPN service, and want to exchange routing information between their sites using OSPF (during migration or on a permanent basis) over an MPLS VPN backbone.

Figure 1 shows an example of how VPN client sites that run OSPF can connect over an MPLS VPN backbone.

Figure 1 *OSPF Connectivity Between VPN Client Sites and an MPLS VPN Backbone*



When OSPF is used to connect PE and CE routers, all routing information learned from a VPN site is placed in the VPN routing and forwarding (VRF) instance associated with the incoming interface. The PE routers that attach to the VPN use the Border Gateway Protocol (BGP) to distribute VPN routes to each other. A CE router can then learn the routes to other sites in the VPN by peering with its attached PE router. The MPLS VPN superbackbone provides an additional level of routing hierarchy to interconnect the VPN sites running OSPF.

When OSPF routes are propagated over the MPLS VPN backbone, additional information about the prefix in the form of BGP extended communities (route type, domain ID extended communities) is appended to the BGP update. This community information is used by the receiving PE router to decide the type of link-state advertisement (LSA) to be generated when the BGP route is redistributed to the OSPF PE-CE process. In this way, internal OSPF routes that belong to the same VPN and are advertised over the VPN backbone are seen as interarea routes on the remote sites.

For basic information about how to configure an MPLS VPN, refer to:


<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t5/vpn.htm>

70390

Using a Sham-Link to Correct OSPF Backdoor Routing

Although OSPF PE-CE connections assume that the only path between two client sites is across the MPLS VPN backbone, backdoor paths between VPN sites (shown in grey in Figure 2) may exist. If these sites belong to the same OSPF area, the path over a backdoor link will always be selected because OSPF prefers intra-area paths to interarea paths. (PE routers advertise OSPF routes learned over the VPN backbone as interarea paths.) For this reason, OSPF backdoor links between VPN sites must be taken into account so that routing is performed based on policy.

Figure 2 Backdoor Paths Between OSPF Client Sites



For example, Figure 2 shows three client sites, each with backdoor links. Because each site runs OSPF within the same Area 1 configuration, all routing between the three sites follows the intra-area path across the backdoor links, rather than over the MPLS VPN backbone.

The following example shows BGP routing table entries for the prefix 10.3.1.7/32 in the PE-1 router in Figure 2. This prefix is the loopback interface of the Winchester CE router. As shown in bold in this example, the loopback interface is learned via BGP from PE-2 and PE-3. It is also generated through redistribution into BGP on PE-1.

```
PE-1# show ip bgp vpnv4 all 10.2.1.7
BGP routing table entry for 100:251:10.3.1.7/32, version 58
Paths: (3 available, best #2)
  Advertised to non peer-group peers:
    10.3.1.2 10.3.1.5
  Local
    10.3.1.5 (metric 30) from 10.3.1.5 (10.3.1.5)
      Origin incomplete, metric 22, localpref 100, valid, internal
      Extended Community: RT:1:793 OSPF DOMAIN ID:0.0.0.100 OSPF
      RT:1:2:0 OSPF 2
  Local
    10.2.1.38 from 0.0.0.0 (10.3.1.6)
      Origin incomplete, metric 86, localpref 100, weight 32768,
      valid, sourced, best
```

```

Extended Community: RT:1:793 OSPF DOMAIN ID:0.0.0.100 OSPF
RT:1:2:0 OSPF 2
Local
10.3.1.2 (metric 30) from 10.3.1.2 (10.3.1.2)
Origin incomplete, metric 11, localpref 100, valid, internal
Extended Community: RT:1:793 OSPF DOMAIN ID:0.0.0.100 OSPF
RT:1:2:0 OSPF 2

```

Within BGP, the locally generated route (10.2.1.38) is considered to be the best route. However, as shown in bold in the next example, the VRF routing table shows that the selected path is learned via OSPF with a next hop of 10.2.1.38, which is the Vienna CE router.

```

PE-1# show ip route vrf ospf 10.3.1.7
Routing entry for 10.3.1.7/32
  Known via "ospf 100", distance 110, metric 86, type intra area
  Redistributing via bgp 215
  Advertised by bgp 215
  Last update from 10.2.1.38 on Serial0/0/0, 00:00:17 ago
  Routing Descriptor Blocks:
  * 10.2.1.38, from 10.3.1.7, 00:00:17 ago, via Serial0/0/0
    Route metric is 86, traffic share count is 1

```


This path is selected because:

- The OSPF intra-area path is preferred over the interarea path (over the MPLS VPN backbone) generated by the PE-1 router.
- OSPF has a lower administrative distance (AD) than internal BGP (BGP running between routers in the same autonomous system).

If the backdoor links between sites are used only for backup purposes and do not participate in the VPN service, then the default route selection shown in the preceding example is not acceptable. To reestablish the desired path selection over the MPLS VPN backbone, you must create an additional OSPF intra-area (logical) link between ingress and egress VRFs on the relevant PE routers. This link is called a sham-link.

A sham-link is required between any two VPN sites that belong to the same OSPF area and share an OSPF backdoor link. If no backdoor link exists between the sites, no sham-link is required.

Figure 3 shows a sample sham-link between PE-1 and PE-2. A cost is configured with each sham-link and is used to decide whether traffic will be sent over the backdoor path or the sham-link path. When a sham-link is configured between PE routers, the PEs can populate the VRF routing table with the OSPF routes learned over the sham-link.

Figure 3 Using a Sham-Link Between PE Routers to Connect OSPF Client Sites

Because the sham-link is seen as an intra-area link between PE routers, an OSPF adjacency is created and database exchange (for the particular OSPF process) occurs across the link. The PE router can then flood LSAs between sites from across the MPLS VPN backbone. As a result, the desired intra-area connectivity is created.


The section, “[Creating a Sham-Link](#)”, describes how to configure a sham-link between two PE routers. For more information about how to configure OSPF, refer to:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cp1/1cospf.htm

Sham-Link Configuration Example

The example in this section is designed to show how a sham-link is used only to affect the OSPF intra-area path selection of the PE and CE routers. The PE router also uses the information received from MP-BGP to set the outgoing label stack of incoming packets, and to decide to which egress PE router to label switch the packets.

[Figure 4](#) shows a sample MPLS VPN topology in which a sham-link configuration is necessary. A VPN client has three sites, each with a backdoor link. Two sham-links have been configured, one between PE-1 and PE-2, and another between PE-2 and PE-3. A sham-link between PE-1 and PE-3 is not necessary in this configuration because the Vienna and Winchester sites do not share a backdoor link.

Figure 4 Sham-Link Example

The following example shows the forwarding that occurs between sites from the standpoint of how PE-1 views the 10.3.1.7/32 prefix, the loopback1 interface of the Winchester CE router in [Figure 4](#).

```
PE-1# show ip bgp vpnv4 all 10.3.1.7
BGP routing table entry for 100:251:10.3.1.7/32, version 124
Paths: (1 available, best #1)
  Local
    10.3.1.2 (metric 30) from 10.3.1.2 (10.3.1.2)
      Origin incomplete, metric 11, localpref 100, valid, internal,
      best
      Extended Community: RT:1:793 OSPF DOMAIN ID:0.0.0.100 OSPF
      RT:1:2:0 OSPF 2

PE-1# show ip route vrf ospf 10.3.1.7
Routing entry for 10.3.1.7/32
  Known via "ospf 100", distance 110, metric 13, type intra area
  Redistributing via bgp 215
  Last update from 10.3.1.2 00:12:59 ago
  Routing Descriptor Blocks:
  10.3.1.2 (Default-IP-Routing-Table), from 10.3.1.7, 00:12:59 ago
```

The next example shows forwarding information in which the next hop for the route, 10.3.1.2, is the PE-3 router rather than the PE-2 router (which is the best path according to OSPF). The reason the OSPF route is not redistributed to BGP on the PE is because the other end of the sham-link already redistributed the route to BGP and there is no need for duplication. The OSPF sham-link is used only to influence intra-area path selection. When sending traffic to a particular destination, the PE router uses the MP-BGP forwarding information.

```
PE-1# show ip bgp vpnv4 all tag | begin 10.3.1.7
10.3.1.7/32      10.3.1.2      notag/38

PE-1# show tag-switching forwarding 10.3.1.2
Local  Outgoing  Prefix  Bytes tag  Outgoing  Next Hop
tag    tag or VC  or Tunnel Id  switched  interface
31     42        10.3.1.2/32    0         PO3/0/0    point2point
```

```

PE-1# show ip cef vrf ospf 10.3.1.7
10.3.1.7/32, version 73, epoch 0, cached adjacency to POS3/0/0
0 packets, 0 bytes
tag information set
  local tag: VPN-route-head
  fast tag rewrite with POS3/0/0, point2point, tags imposed: {42 38}
via 10.3.1.2, 0 dependencies, recursive
  next hop 10.1.1.17, POS3/0/0 via 10.3.1.2/32
  valid cached adjacency
  tag rewrite with POS3/0/0, point2point, tags imposed: {42 38}

```

If a prefix is learned across the sham-link and the path via the sham-link is selected as the best, the PE router does not generate an MP-BGP update for the prefix. It is not possible to route traffic from one sham-link over another sham-link.

In the following example, PE-2 shows how an MP-BGP update for the prefix is not generated. Although 10.3.1.7/32 has been learned via OSPF across the sham-link as shown in bold, no local generation of a route into BGP is performed. The only entry within the BGP table is the MP-BGP update received from PE-3 (the egress PE router for the 10.3.1.7/32 prefix).

```

PE-2# show ip route vrf ospf 10.3.1.7
Routing entry for 10.3.1.7/32
  Known via "ospf 100", distance 110, metric 12, type intra area
  Redistributing via bgp 215
  Last update from 10.3.1.2 00:00:10 ago
  Routing Descriptor Blocks:
    * 10.3.1.2 (Default-IP-Routing-Table), from 10.3.1.7, 00:00:10 ago
      Route metric is 12, traffic share count is 1

PE-2# show ip bgp vpnv4 all 10.3.1.7
BGP routing table entry for 100:251:10.3.1.7/32, version 166
Paths: (1 available, best #1)
  Not advertised to any peer
  Local
    10.3.1.2 (metric 30) from 10.3.1.2 (10.3.1.2)
      Origin incomplete, metric 11, localpref 100, valid, internal,
      best
      Extended Community: RT:1:793 OSPF DOMAIN ID:0.0.0.100 OSPF
      RT:1:2:0 OSPF 2

```

The PE router uses the information received from MP-BGP to set the ongoing label stack of incoming packets, and to decide to which egress PE router to label switch the packets.

Benefits

Client Site Connection Across the MPLS VPN Backbone

A sham-link overcomes the OSPF default behavior for selecting an intra-area backdoor route between VPN sites instead of an interarea (PE-to-PE) route. A sham-link ensures that OSPF client sites that share a backdoor link can communicate over the MPLS VPN backbone and participate in VPN services.

Flexible Routing in an MPLS VPN Configuration

In an MPLS VPN configuration, the OSPF cost configured with a sham-link allows you to decide if OSPF client site traffic will be routed over a backdoor link or through the VPN backbone.

Restrictions

When OSPF is used as a protocol between PE and CE routers, the OSPF metric is preserved when routes are advertised over the VPN backbone. The metric is used on the remote PE routers to select the correct route. For this reason, you should not modify the metric value when OSPF is redistributed to BGP, and when BGP is redistributed to OSPF. If you modify the metric value, routing loops may occur.

Related Features and Technologies

- MPLS
- OSPF
- BGP

Related Documents

- *Cisco IOS Configuration Fundamentals Configuration Guide, Release 12.0*
- *Cisco IOS Configuration Fundamentals Command Reference, Release 12.0*
- *Network Protocols Command Reference, Part 1, Release 11.3*
http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/113ed_cr/np1_r/1rospf.htm
- *MPLS Virtual Private Networks*
<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t5/vpn.htm>
- *Configuring OSPF*
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cp1/1cospf.htm#6844
- *Configuring BGP*
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cp1/1cbgp.htm
- RFC 1163, A Border Gateway Protocol
- RFC 1164, Application of the Border Gateway Protocol in the Internet
- RFC 2283, Multiprotocol Extensions for BGP-4
- RFC 2328, Open Shortest Path First, Version 2
- RFC 2547, BGP/MPLS VPNs

Supported Platforms

Determining Platform Support Through Cisco Feature Navigator

Cisco IOS software is packaged in feature sets that support specific platforms. To get updated information regarding platform support for this feature, access Cisco Feature Navigator. Cisco Feature Navigator dynamically updates the list of supported platforms as new platform support is added for the feature.

Cisco Feature Navigator is a web-based tool that enables you to quickly determine which Cisco IOS software images support a specific set of features and which features are supported in a specific Cisco IOS image. You can search by feature or release. Under the release section, you can compare releases side by side to display both the features unique to each software release and the features in common.

To access Cisco Feature Navigator, you must have an account on Cisco.com. If you have forgotten or lost your account information, send a blank e-mail to cco-locksmith@cisco.com. An automatic check will verify that your e-mail address is registered with Cisco.com. If the check is successful, account details with a new random password will be e-mailed to you. Qualified users can establish an account on Cisco.com by following the directions at <http://www.cisco.com/register>.

Cisco Feature Navigator is updated regularly when major Cisco IOS software releases and technology releases occur. For the most current information, go to the Cisco Feature Navigator home page at the following URL:

<http://www.cisco.com/go/fn>

Supported Standards, MIBs, and RFCs

Standards

No new or modified standards are supported by this feature.

MIBs

No new or modified MIBs are supported by this feature.

To obtain lists of supported MIBs by platform and Cisco IOS release, and to download MIB modules, go to the Cisco MIB website on Cisco.com at the following URL:

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

RFCs

No new or modified RFCs are supported by this feature.

Prerequisites

Before you can configure a sham-link in an MPLS VPN, you must first enable OSPF as follows:

- Create an OSPF routing process.
- Specify the range of IP addresses to be associated with the routing process.
- Assign area IDs to be associated with the range of IP addresses.

For more information on these OSPF configuration procedures, go to:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cprt1/1cospf.htm#6844

Configuration Tasks

See the following sections for configuration tasks for the sham-link feature. Each task in the list is identified as either required or optional.

- [Creating a Sham-Link](#) (required)
- [Verifying Sham-Link Creation](#) (optional)

Creating a Sham-Link

Before you create a sham-link between PE routers in an MPLS VPN, you must:

- Configure a separate /32 address on the remote PE so that OSPF packets can be sent over the VPN backbone to the remote end of the sham-link. The /32 address must meet the following criteria:
 - Belong to a VRF.
 - Not be advertised by OSPF.
 - Be advertised by BGP.

You can use the /32 address for other sham-links.

- Associate the sham-link with an existing OSPF area.

To create a sham-link, use the following commands starting in EXEC mode:

	Command	Purpose
Step 1	Router1# configure terminal	Enters global configuration mode on the first PE router.
Step 2	Router1(config)# interface loopback <i>interface-number</i>	Creates a loopback interface to be used as an endpoint of the sham-link on PE-1 and enters interface configuration mode.
Step 3	Router1(config-if)# ip vrf forwarding <i>vrf-name</i>	Associates the loopback interface with a VRF. Removes the IP address.
Step 4	Router1(config-if)# ip address <i>ip-address</i> <i>mask</i>	Reconfigures the IP address of the loopback interface on PE-1.
Step 5	Router1(config-if)# end	Returns to global configuration mode.
Step 6	Router1(config)# end	Returns to EXEC mode.
Step 7	Router2# configure terminal	Enters global configuration mode on the second PE router.
Step 8	Router2(config)# interface loopback <i>interface-number</i>	Creates a loopback interface to be used as the endpoint of the sham-link on PE-2 and enters interface configuration mode.
Step 9	Router2(config-if)# ip vrf forwarding <i>vrf-name</i>	Associates the second loopback interface with a VRF. Removes the IP address.
Step 10	Router2(config-if)# ip address <i>ip-address</i> <i>mask</i>	Reconfigures the IP address of the loopback interface on PE-2.
Step 11	Router2(config-if)# end	Returns to global configuration mode.
Step 12	Router1(config)# end	Returns to EXEC mode.
Step 13	Router1(config)# router ospf <i>process-id</i> vrf <i>vrf-name</i>	Configures the specified OSPF process with the VRF associated with the sham-link interface on PE-1 and enters interface configuration mode.
Step 14	Router1(config-if)# area <i>area-id</i> sham-link <i>source-address</i> <i>destination-address</i> cost <i>number</i>	Configures the sham-link on the PE-1 interface within a specified OSPF area and with the loopback interfaces specified by the IP addresses as endpoints. cost number configures the OSPF cost for sending an IP packet on the PE-1 sham-link interface.

	Command	Purpose
Step 15	Router2(config)# router ospf <i>process-id</i> vrf <i>vrf-name</i>	Configures the specified OSPF process with the VRF associated with the sham-link interface on PE-2 and enters interface configuration mode.
Step 16	Router2(config-if)# area <i>area-id</i> sham-link <i>source-address</i> <i>destination-address</i> cost <i>number</i>	Configures the sham-link on the PE-2 interface within a specified OSPF area and with the loopback interfaces specified by the IP addresses as endpoints. cost number configures the OSPF cost for sending an IP packet on the PE-2 sham-link interface.

Verifying Sham-Link Creation

To verify that the sham-link was successfully created and is operational, use the **show ip ospf sham-links** command in EXEC mode:

```
Router1# show ip ospf sham-links

Sham Link OSPF_SL0 to address 10.2.1.2 is up
Area 1 source address 10.2.1.1
Run as demand circuit
DoNotAge LSA allowed. Cost of using 40 State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40,
Hello due in 00:00:04
Adjacency State FULL (Hello suppressed)
Index 2/2, retransmission queue length 4, number of
retransmission 0
First 0x63311F3C(205)/0x63311FE4(59) Next
0x63311F3C(205)/0x63311FE4(59)
Last retransmission scan length is 0, maximum is 0
Last retransmission scan time is 0 msec, maximum is 0 msec
Link State retransmission due in 360 msec
```

Monitoring and Maintaining a Sham-Link

To monitor a sham-link, use the following **show** commands in EXEC mode:

Command	Purpose
Router# show ip ospf sham-links	Displays the operational status of all sham-links configured for a router.
Router# show ip ospf data router <i>ip-address</i>	Displays information about how the sham-link is advertised as an unnumbered point-to-point connection between two PE routers.

Configuration Examples

The following example shows how to configure a sham-link between two PE routers:

```
Router1(config)# interface loopback 1
Router1(config-if)# ip vrf forwarding ospf
Router1(config-if)# ip address 10.2.1.1 255.255.255.255
!
Router2(config)# interface loopback 1
```

```
Router2(config-if)# ip vrf forwarding ospf
Router2(config-if)# ip address 10.2.1.2 255.255.255.255
!
Router1(config)# router ospf 100 vrf ospf
Router1(config-if)# area 1 sham-link 10.2.1.1 10.2.1.2 cost 40
!
Router2(config)# router ospf 100 vrf ospf
Router2(config-if)# area 1 sham-link 10.2.1.2 10.2.1.1 cost 40
```

Command Reference

The following commands are introduced or modified in the feature or features documented in this module. For information about these commands, see the *Cisco IOS Multiprotocol Label Switching Command Reference* at http://www.cisco.com/en/US/docs/ios/mpls/command/reference/mp_book.html. For information about all Cisco IOS commands, go to the Command Lookup Tool at <http://tools.cisco.com/Support/CLILookup> or to the *Cisco IOS Master Commands List*.

- **area sham-link cost**
- **show ip ospf sham-links**

Glossary

BGP—Border Gateway Protocol. Interdomain routing protocol that exchanges reachability information with other BGP systems. It is defined in RFC 1163.

CE router—customer edge router. A router that is part of a customer network and that interfaces to a provider edge (PE) router. CE routers are not aware of associated VPNs.

CEF—Cisco Express Forwarding. An advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with large and dynamic traffic patterns.

OSPF—Open Shortest Path First protocol.

IGP—Interior Gateway Protocol. An Internet protocol used to exchange routing information within an autonomous system. Examples of common IGP include IGRP, OSPF, and RIP.

LSA—link-state advertisement. A broadcast packet used by link-state protocols. The LSA contains information about neighbors and path costs and is used by the receiving router to maintain a routing table.

MPLS—Multiprotocol Label Switching. Emerging industry standard upon which tag switching is based.

PE router—provider edge router. A router that is part of a service provider network connected to a customer edge (CE) router. All VPN processing occurs in the PE router.

SPF—shortest path first calculation.

VPN—Virtual Private Network. A secure IP-based network that shares resources on one or more physical networks. A VPN contains geographically dispersed sites that can communicate securely over a shared backbone.

VRF—VPN routing and forwarding instance. A VRF consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table. In general, a VRF includes the routing information that defines a customer VPN site that is attached to a PE router.

CCVP, the Cisco logo, and Welcome to the Human Network are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0711R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2007 Cisco Systems, Inc. All rights reserved.

