



Cisco MDS 9000 Family FCIP WAN Deployment Guidelines

SAN-OS 1.1(2)

Introduction

FCIP, a transparent Fibre Channel tunneling protocol over TCP/IP, can be used for a variety of applications across a variety of network transports and topologies. FCIP can be used across both high and low speed links, and across long and short distances and latencies.

When considering various transports of differing speed and distance, it is important to tune the FCIP transport to ensure expected performance and resiliency. For example, special configuration precautions must be followed when building FCIP solutions over very long, slow WAN facilities.

The following best practices and design guidelines are presented to help the SAN designer to build FCIP-based solutions that are optimized in terms of performance and resiliency. The following design guidelines are based on laboratory testing at Cisco Systems.

Applicable SAN-OS

The following configuration guidelines are based upon MDS9000 SAN-OS system and kickstart version 1.1(2) and higher.

This release introduces TCP packet shaping to FCIP. TCP packet shaping enables the sender to immediately send packets at the configured minimum available bandwidth of the FCIP path.

FCIP Configuration

The following guidelines assume a *dedicated* FCIP WAN link where the minimum available bandwidth is equal to the maximum available bandwidth (or the speed of the WAN link).

Note that whilst `tcp cwm burstsize` parameter has no effect when `max-bandwidth-kbps = min-available-bandwidth-kbps`, it should remain enabled to prevent large bursts at gigabit rate. i.e the shaper is not enabled after partial idle periods unless cwm is enabled.

Guideline configuration for an FCIP entity is shown below.

Recommended Configuration for FCIP over dedicated WAN link

```
fcip profile 10
ip address 10.1.1.1
tcp sack-enable
tcp cwm burstsize 20
tcp max-bandwidth-kbps <path bandwidth> min-available-bandwidth-kbps
<path bandwidth> round-trip-time-ms <RTT>
```



```
interface fcipl
switchport trunk mode off
no shutdown
no switchport trunk allowed vsan all
switchport trunk allowed vsan add 5
use-profile 10
peer-info ipaddr 10.1.1.4

interface GigabitEthernet2/1
ip address 10.1.1.1 255.255.255.0
switchport mtu 2500
no shutdown
```

Notes:

1. The MTU size must be configured to no more than 2500 (this will still hold a maximum size Fibre Channel frame, so anything larger is superfluous anyway).
2. SACK (Selective ACK) must be enabled at each end of the FCIP link (SACK is enabled by default with this release).

Examples:

1. Gigabit Ethernet link at 20ms Round trip delay
`tcp max-bandwidth-mbps 1000 min-available-bandwidth-mbps 1000 round-trip-time-ms 20`
2. OC-12 link at 10ms Round trip delay
`tcp max-bandwidth-mbps 622 min-available-bandwidth-mbps 622 round-trip-time-ms 10`
3. OC-3 link at 10ms Round trip delay
`tcp max-bandwidth-mbps 155 min-available-bandwidth-mbps 155 round-trip-time-ms 10`

Guidelines for Maximum Fibre Channel Sources at GigE, OC-12, OC-3, and T3 Rates

The following are guidelines based upon laboratory testing for the maximum number of Fibre Channel sources for given link bandwidths and round trip times (RTT).

The following guidelines cover Gigabit Ethernet (1000Mbps), OC-12 (622Mbps), OC-3 (155Mbps), and T3 (45 Mbps) WAN rates at round trip times (RTT) of up to 100ms (50ms each way) and IP WAN packet loss of 0, 0.01%, 0.05%, 0.1%, 0.5% and 1% (random packet loss).

Note that the guideline matrices for packet loss are built around random statistical distributions for packet loss.

For each matrix, the left column shows the number of FC sources or flows, each at 100 MBps (1 Gbps). The top row indicates the RTT (Round trip time) in milliseconds. The intersection of the RTT and flows give the suggested number of Receive BB Credits for ports destined over the FCIP WAN link. Receive BB Credits can be configured through the Device Manager or through the CLI (`switchport forxbbcredit <xx>`). The default number of BB credits allocated to an F_Port or FL_Port is 16 on the Cisco MDS 9000 Family switches and therefore may not require any changes for the following examples. Invalid configurations (configurations that may lead to Fibre Channel frame expiry) are denoted by the shaded boxes labeled "Not Valid".



Gigabit Ethernet

Notes: Each flow generated at 100MBps with 2148 Byte frames (one way)

Packet loss: Zero

One way FC throughput: ~115MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.01%

One way FC throughput: ~77MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.05%

One way FC throughput: ~50MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB



Flows	10ms	20ms	50ms	100ms
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.1%

One way FC throughput: ~42MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.5%

One way FC throughput: ~25MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB



Packet loss: 1%

One way FC throughput: ~17MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

OC-12 (622Mbps)

Notes: Each flow generated at 100MBps with 2148 Byte frames (one way)

Packet loss: Zero

One way FC throughput: ~71MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.01%

One way FC throughput: ~53MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB



Flows	10ms	20ms	50ms	100ms
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.05%

One way FC throughput: ~36MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.1%

One way FC throughput: ~32MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB



Packet loss: 0.5%

One way FC throughput: ~20MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 1%

One way FC throughput: ~14MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

OC-3 (155Mbps)

Notes: Each flow generated at 100MBps with 2148 Byte frames (one way)

Packet loss: Zero

One way FC throughput: ~17.8MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB



Flows	10ms	20ms	50ms	100ms
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.01%

One way FC throughput: ~15.6MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.05%

One way FC throughput: ~12.8MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB



Packet loss: 0.1%

One way FC throughput: ~11MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.5%

One way FC throughput: ~7.6MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 1%

One way FC throughput: ~6.7MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB



Flows	10ms	20ms	50ms	100ms
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

T3 (45Mbps)

Notes: Each flow generated at 100MBps with 2148 Byte frames (one way)

Packet loss: Zero

One way FC throughput: ~5.2MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.01%

One way FC throughput: ~5.0MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB



Packet loss: 0.05%

One way FC throughput: ~4.5MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB

Packet loss: 0.1%

One way FC throughput: ~4.1MBps (@100ms RTT)

Flows	10ms	20ms	50ms	100ms
8	16 BB	16 BB	16 BB	16 BB
7	16 BB	16 BB	16 BB	16 BB
6	16 BB	16 BB	16 BB	16 BB
5	16 BB	16 BB	16 BB	16 BB
4	16 BB	16 BB	16 BB	16 BB
3	16 BB	16 BB	16 BB	16 BB
2	16 BB	16 BB	16 BB	16 BB
1	16 BB	16 BB	16 BB	16 BB



Packet loss: 0.5%

One way FC throughput: ~3.0MBps (@100ms RTT)

Alter default Fibre Channel Receive Performance Buffers for high RTT and # sources as follows for each ingress fc interface:

```
switchport fcrxbbcredit performance-buffers 4
```

To return performance buffers to default value, issue:

```
switchport fcrxbbcredit performance-buffers default
```

Flows	10ms	20ms	50ms	75ms	100ms
8	2 BB	2 BB	2 BB	2 BB	2 BB 4 PerfBfr
7	2 BB	2 BB	2 BB	2 BB	2 BB 4 PerfBfr
6	2 BB	2 BB	2 BB	2 BB	2 BB
5	2 BB	2 BB	2 BB	2 BB	2 BB
4	2 BB	2 BB	2 BB	2 BB	2 BB
3	2 BB	2 BB	2 BB	2 BB	2 BB
2	2 BB	2 BB	2 BB	2 BB	2 BB
1	2 BB	2 BB	2 BB	2 BB	2 BB

Packet loss: 1%

One way FC throughput: ~2.5MBps (@100ms RTT)

Alter default Fibre Channel Receive Performance Buffers for high RTT and # sources as follows for each ingress fc interface:

```
switchport fcrxbbcredit performance-buffers 4
```

To return performance buffers to default value, issue:

```
switchport fcrxbbcredit performance-buffers default
```

Flows	10ms	20ms	50ms	75ms	100ms
8	2 BB	2 BB	2 BB	2 BB 4 PerfBfr	2 BB 4 PerfBfr
7	2 BB	2 BB	2 BB	2 BB 4 PerfBfr	2 BB 4 PerfBfr
6	2 BB	2 BB	2 BB	2 BB 4 PerfBfr	2 BB 4 PerfBfr
5	2 BB	2 BB	2 BB	2 BB 4 PerfBfr	2 BB 4 PerfBfr
4	2 BB	2 BB	2 BB	2 BB	2 BB
3	2 BB	2 BB	2 BB	2 BB	2 BB
2	2 BB	2 BB	2 BB	2 BB	2 BB
1	2 BB	2 BB	2 BB	2 BB	2 BB



Latency through FCIP Tunnel

Fibre Channel frames are buffered in each switch as they move through the network. The length of time each frame is buffered is a function of the buffer depth, and the rate at which that buffer can be emptied. In the case of FCIP links, the process to empty the buffer can be slower according to the available bandwidth of the WAN link supporting the FCIP tunnel and the presence of any packet drops and retransmissions over that link.

End-to-end Latency at low traffic rates

At low traffic rates where no queuing occurs from traffic congestion, end-to-end latency for a 2148 Byte Fibre Channel frame through the FCIP tunnel is around 220 including serialization time through two MDS9000 switches in a back-to-back configuration. Transit delay, serialization, and store-and-forward delays through intermediate routers and switches in the IP network are additional.

End-to-end latency at high traffic rates

Higher latencies will occur whenever queuing occurs in the network. The upper bound for latency is reached when the ingress rate exceeds the capacity of the FCIP link. When this occurs, the Fibre Channel Receive Buffers are filled, thus withholding the return of R-Rdys to the source(s) nodes.

Hence an approximate upper bound for latency can be calculated as follows:

$$L = (\#Bfrs / \text{Send Rate for FCIP link}) + \text{end-to-end latency}$$

For example,

45Mbps FCIP link with 8x 1Gbps FC sources each with 16 Rx BB Credits and Performance Buffers at default. Frame Size is 2kB average. End-to-end latency is 50ms.

$$45\text{Mbps link} \Rightarrow 45E6 / 8 / 2048 = 2746 \text{ packets per second send rate for FCIP link}$$

$$8x \text{ 1Gbps FC sources} \Rightarrow 8 \times (16+16) \text{ buffers} = 256 \text{ receive buffers}$$

$$L = (256 / 2746) + 0.050 = 0.093 + 0.050 = 0.143\text{seconds} = 143 \text{ ms one way latency upper bound (approximate)}$$

Effect of Various Parameters on Overall Latency

As can be seen from the example above, when congestion occurs, larger Fibre Channel receive buffers will contribute to a larger network latency. However, since the alternative is to push back on the data source, the overall effect is arguably negligible—the frames are buffered either in the switch or at the source.

Naturally, a larger available bandwidth for the FCIP link will improve this situation by effectively raising the “congestion bar.” However, queuing and congestion will always occur when the ingress data sources(s) exceed the capacity of the FCIP link.



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2003 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site 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.
(0304R) ETMG 203154—LSK 09/03