



NETWORKERS 2004

HOW CISCO ACHIEVED HIGH AVAILABILITY IN ITS LOCAL AREA NETWORK

SESSION NMS-2202

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

1

Agenda

Cisco.com

- **How Cisco IT Measures Availability**
- **Cisco IT-LAN San Jose CY '03 Availability Results**
- **Best Practices for High Availability**
- **Cisco Specific Considerations**
- **High-Availability Technology Choices**

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

2

Reference Materials

Cisco.com

- **Cisco High Availability White Papers**

http://www.cisco.com/en/US/tech/tk869/tk769/tech_white_papers_list.html

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

3

Associated Sessions

Cisco.com

- **NMS-2201: Network Availability Measurement**
- **NMS-2T20: Designing and Managing High Availability IP Networks**
- **RST-2514: High Availability in Campus Network Deployments**
- **RST-4312: High Availability in Routing**

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

4

HOW CISCO IT MEASURES AVAILABILITY



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

5

The Cisco San Jose LAN/MAN

Cisco.com

- 50 buildings
- 800 access-layer switches
- 900 wireless access points
- 2 production datacenters
- 4 development datacenters



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

6

How Cisco IT Measures Raw Availability

Cisco.com

- **Availability should estimate the client experience**
- **Cisco IT uses the uptime of the access layer switch to estimate the customer's uptime**
- **Cisco pings every access layer switch twice every 15–20 seconds**
 - If either ping responds the device is considered “up”**
 - If neither ping responds the device is considered “down”**

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

7

Change Management Procedure

Cisco.com

- **Network engineers fill out a web form prior to a change**
 - Who/what/where/when/why?**
 - What devices will be down? When will they be down?**
- **The change process includes notifying affected customers (usually via email sent to a building)**
- **Results:**
 - Imposes discipline during change planning**
 - Notifications help to minimize “change collision”**
 - Provides a record of changes which facilitates problem resolution**
 - Operations Command Center knows whether outages in progress are planned or unplanned**


NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

8

Raw vs. Adjusted Availability

Cisco.com



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

9

- Our change management system is integrated with our availability measurement system

We know what devices are planned to be down
We know when they are planned to be down

- This allows us to calculate both “raw” and “adjusted” availability

“Adjusted” availability considers devices 100% up during an approved change

“Raw” availability ignores changes

- Both metrics are valuable

Aggregating Availability Statistics

Cisco.com

- Devices are organized into subgroups (examples: production datacenter 1; site 4 desktops)
- Device availability is averaged to calculate the subgroup availability
- All SJ LAN subgroups are in the SJ_LAN_MAN group
- SJ_LAN_MAN availability is the average of all devices (not subgroups!)

- Group: **SJ_LAN_MAN** - San Jose LAN and MAN
 - SubGroup: **DDC1** - Development Datacenter 1
 - SubGroup: **DDC2** - Development Datacenter 2
 - SubGroup: **DDC3** - Development Datacenter 3
 - SubGroup: **DDC4** - Development Datacenter 4
 - SubGroup: **Lab** - Lab Networks
 - SubGroup: **MAN** - MAN Networks
 - SubGroup: **PDC1** - Production Datacenter 1
 - SubGroup: **PDC2** - Production Datacenter 2
 - SubGroup: **Site 1** - Site 1-3 Desktop Networks
 - SubGroup: **Site 4** - Site 4 Desktop Networks
 - SubGroup: **Site 5** - Site 5 Desktop Networks

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

10

Availability throughout Cisco

Cisco.com

- **Network teams ping their devices**
- **Hosting teams ping their servers**
- **Webmaster teams test http against web servers**
- **Application teams test synthetic web and database transactions**

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

11

Is Availability a Perfect Measurement?

Cisco.com

- **Some outages affect customers but do not impact availability**
Example: Rogue DHCP server
- **Linecard failures impact users but do not always make the switch unpingable**
- **Solution:**
Cisco IT Networking reports 3 key metrics in our quarterly operations reviews:
Availability
P1 metrics
Detailed information on any high severity P1's

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

12

Availability Tool Improvements Wanted

Cisco.com

- Time-based availability: P1 devices measured 24x7; P2 devices measured 12x5
- Availability measurements weighed by number of ports in the device

...but don't make the system too labor-intensive to maintain...

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

13

Cisco IT-LAN San Jose CY '03 Availability

Cisco.com

Subgroup	Adjusted Availability	Raw Availability
Datacenters	99.998%	99.997%
Desktops	99.997%	99.964%
MAN	99.992%	99.985%
Lab	99.987%	99.957%
Overall	99.997%	99.972%

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

14

Q AND A ON MEASURING AVAILABILITY



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

15

OPERATIONAL BEST PRACTICES TO ACHIEVE HIGH AVAILABILITY



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

16

Step #1: Measure Availability

Cisco.com

- **Network engineers must know what availability was yesterday, and why**
- **Cisco IT reports availability monthly, quarterly, and annually**
- **Reports include:**
 - Raw and adjusted availability**
 - Five largest unplanned outages**
 - Five largest planned outages**

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

17

Physical and Logical Hierarchy

Cisco.com

- **Core, distribution, and access layer methodology**
- **Hierarchy built into structured cabling**
- **Hierarchy built into physical network topology**
- **Hierarchy built into IP addressing**


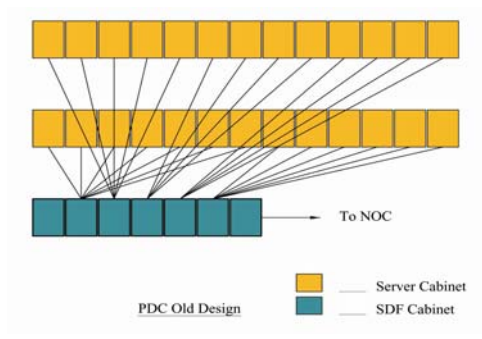
NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

18

Datacenter without Physical Hierarchy

Cisco.com



Source: Cabling Installation and Maintenance, January, 2004

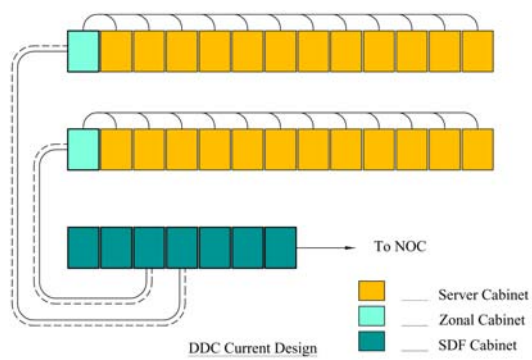
NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

19

Development Datacenter with Physical Hierarchy

Cisco.com



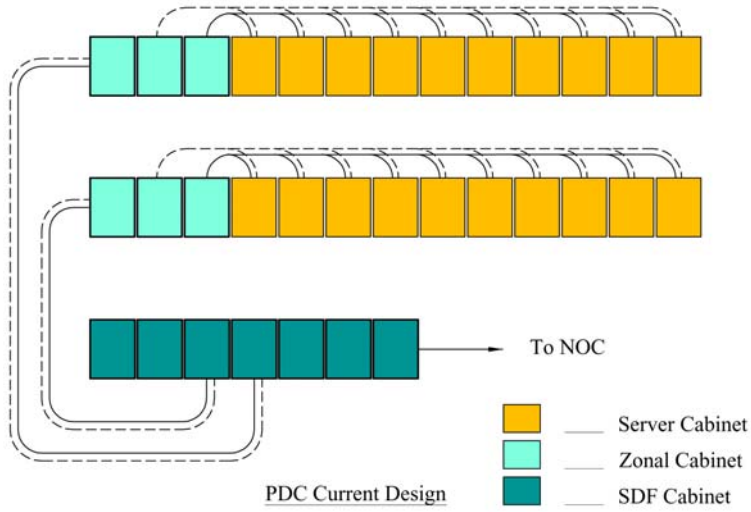
NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

20

Production Datacenter with Physical Hierarchy

Cisco.com




NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

21

Cisco's WAN Network Before Introducing Hierarchy

Cisco.com



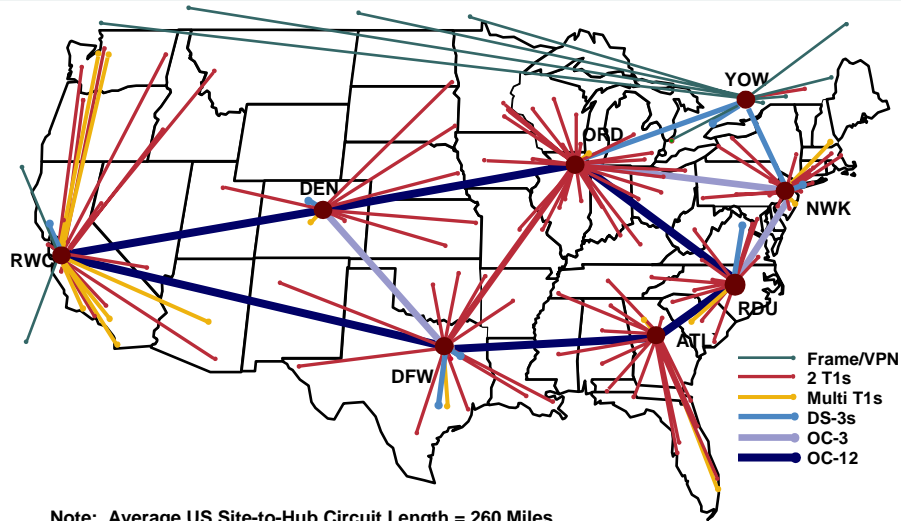
NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

22

Cisco's WAN Network after Introducing Hierarchy

Cisco.com



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

23

Logical Hierarchy

Cisco.com

- Implemented with hierarchical address allocation and EIGRP summarization
- Cisco IT has 23,000 subnets worldwide, but only 800 routes in our core EIGRP routing table
- Every engineer gets an automated daily report on route entry changes in the backbone

Comparing Routes for San Jose RBB

From: Thu Jun 19 7:30:00 US/Pacific 2003

To: Fri Jun 20 7:30:00 US/Pacific 2003

Deleted: 10.75.1.0/30—WAN Link(banaglore-gw1-s3/0)

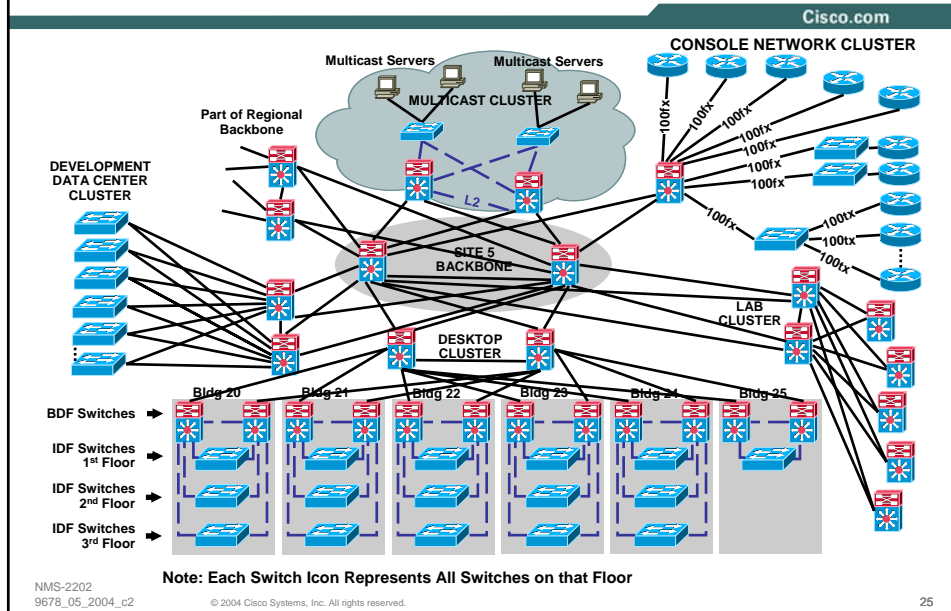
Added: 202.3.11.0/27—

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

24

San Jose Site 5



Outage Root-Cause Analysis

Cisco.com

- Identify the cause of the outage
- Fix the affected building
- Prevent the outage in other buildings
- Update design documentation so new buildings will not be vulnerable
- A uniform network implementation allows lessons learned in one building to benefit all buildings

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

26

Power Considerations

Cisco.com

- **Redundant power supplies should go into separate power circuits (if available)**
- **What parts of the network should be on “house power”?**
- **What parts of the network should be on UPS?**
- **What parts of the network should be on UPS plus generator?**
- **Do you need N+1 generators?**
- **Do you need a refueling contract?**



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

27

Network Redundancy

Cisco.com

- **Provisioning the network with redundancy is only the first step**
- **We use a weekly perl script to verify that each device has two separate paths back to the core**
- **WAN circuit diversity (and LAN fiber path diversity) are constant challenges**
- **Paging for circuit outages helps, but this imposes a high cost on on-call personnel**
- **For critical parts of the network we intentionally test redundancy during planned changes**

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

28

Configuration Audits

Cisco.com

- We use the Router Audit Tool from www.cisecurity.org
- We use a custom rule file for routers, switches, and Cisco IOS® APs
- Functionality mostly limited to global configuration; the tool's ability to understand interface configurations is limited
- We use a Perl script to generate weekly "bottom-10 device" and "10 most violated rules" lists

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

29

Cisco IOS/Catalyst OS Version Standardization

Cisco.com

- Operations teams pass testing and bug information to our global "Cisco IOS version coordinator"
- Coordinator picks "experimental", "recommended", and "acceptable" Cisco IOS versions
- Automated scripts monitor adoption rates and emails exception list to local teams

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

30

What Does 12.1(8b)E10 Mean?

Cisco.com

- **12.0T (technology train) integrated many new features**
- **12.1(1) (mainline, also referred to as 12.1M) was a copy of the last 12.1T release; it is called a “major release”**
- **12.1(8) was the 7th maintenance release of 12.1(1); each maintenance release undergoes full regression testing; new features are rarely introduced into a mainline release**
- **12.1(8b) was the 2nd “rebuild” of 12.1(8); rebuilds include a handpicked set of bug fixes but do not undergo full regression testing**
- **12.1(8b)E integrated “Enterprise Train” features (Catalyst® 6500 for example) with the 12.1(8b) code base**
- **12.1(8b)E10 is the 10th rebuild of 12.1(8b)E**

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

31

Cisco IOS Selection Tactics

Cisco.com

- **Technology (and “early deployment”) trains integrate new features; use them only when you need those new features**
- **Mainline trains rarely introduce new features; mainline trains are either “limited deployment” or “general deployment”**
- **Enterprise and service provider trains introduce a limited number of new features**
- **Pick images with the largest number of rebuilds or maintenance releases since the integration of the latest feature**
- **Some releases go through additional “Safe Harbor” validation testing**

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

32

Separate Incident Response from Outage Management

Cisco.com

Functions of Network On-Call Personnel:

- Fix the technical problem
- Identify root cause
- Prevent outage from reoccurring



Functions of Operations Command Center (OCC):

- Determine the outage severity level
- Escalate based on the outage duration and severity
- Contact additional resources (including vendor support) as necessary
- Communicate with affected users

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

33

Troubleshooting Training

Cisco.com

- Ironically, when achieving 99.99% availability, there are not enough “unplanned outages” to keep troubleshooting skills sharp
- Cisco IT LAN team members deliberately break our lab networks and then other team members fix it (as a training exercise)
- Having a test network with a similar topology to the production network is essential



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

34

Other Items to Remember

Cisco.com

- **Outage alerts**
- **Emergency spares**
- **Out-of-band management**

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

35

The Final “Operational Best Practice”: **Paranoia**

Cisco.com

- **Striving for “five nines” is very challenging**
- **Cisco IT continually asks:**
 - What was availability yesterday?**
 - Why?**
- **This is a stressful way to live**
- **We warn candidates during the hiring process**

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

36

Q AND A ON OPERATIONAL BEST PRACTICES



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

37

CISCO SPECIFIC CONSIDERATIONS



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

38

Separation of Production and Alpha Networks

Cisco.com

- We'd love to test last night's Cisco IOS build on the production network...but...
 - People dial 911 on the production network
- In all buildings, the green and yellow jacks patch to the production network
- In many buildings, the orange jack patches to the alpha network
- Engineering VPs "encourage" their staff to patch to the alpha network



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

39

Yes, Cisco IT Calls TAC

Cisco.com

- Cisco IT Networking emphasizes preventing outages
- As a result of our success, our troubleshooting skills are beginning to diminish
- Cisco TAC troubleshoots outages full-time



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

40

HIGH-AVAILABILITY TECHNOLOGY CHOICES



NMS-2202
9678_05_2004_c2

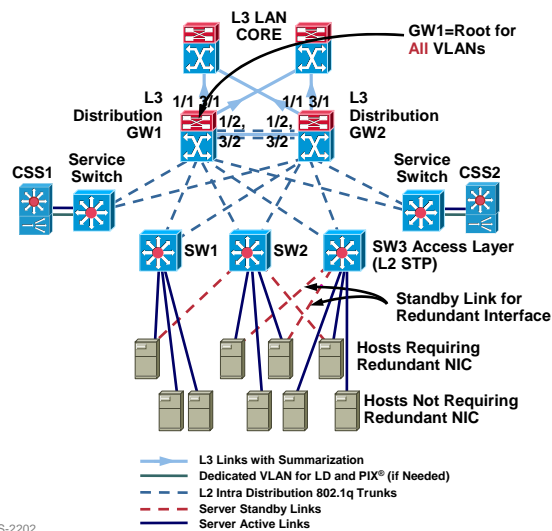
© 2004 Cisco Systems, Inc. All rights reserved.

41

Production Datacenter Design

Cisco.com

STANDARD L2 SOLUTION FOR PDC



NMS-2202
9678_05_2004_c2

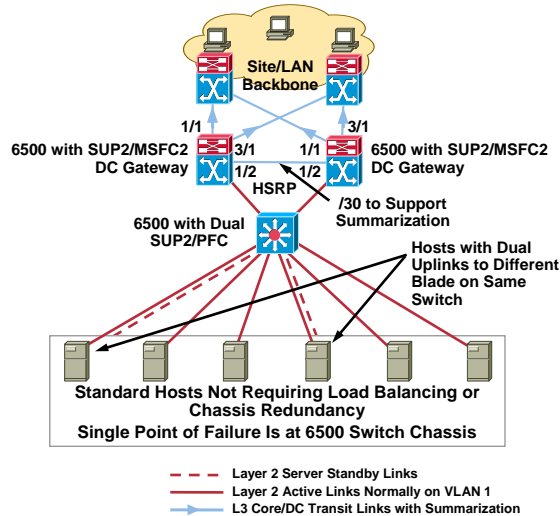
© 2004 Cisco Systems, Inc. All rights reserved.

42

Development Datacenter Design

Cisco.com

STANDARD L3 SOLUTION



NMS-2202
9678_05_2004_c2

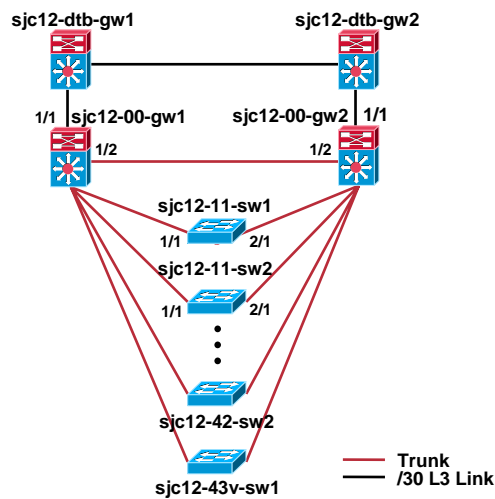
© 2004 Cisco Systems, Inc. All rights reserved.

43

- One subnet/VLAN per switch
- Benefit: No spanning-tree on gateways
- Cost: Does not support fully-redundant services
- GW1 HSRP primary for all VLANs

Production Desktop Network Design

Cisco.com



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

44

- Four VLANs per access-layer switch: Data, voice, wireless data, wireless voice
- SC0 interface on data VLAN to improve availability measurements
- GW1 HSRP primary and spanning-tree root for all VLANs
- Redundant supervisors used in L2 devices
- Single supervisors in L3 devices: EIGRP and HSRP used for redundancy
- VACLs prevent rogue DHCP servers

Q AND A



NMS-2202
9678_05_2004_c2


© 2004 Cisco Systems, Inc. All rights reserved.

45

Recommended Reading

Cisco.com

- **Performance and Fault Management**
ISBN: 1-57870-180-5
- **High Availability Network Fundamentals**
ISBN: 1-57713-017-3
- **Datacenter Fundamentals**
ISBN: 1-57870-023-4



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

46

Complete Your Online Session Evaluation!

Cisco.com

- WHAT:** Complete an online session evaluation and your name will be entered into a daily drawing
- WHY:** Win fabulous prizes! Give us your feedback!
- WHERE:** Go to the Internet stations located throughout the Convention Center
- HOW:** Winners will be posted on the onsite Networkers Website; four winners per day

NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

47

CISCO SYSTEMS



NMS-2202
9678_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

48