



CISCO™

PARTNER

**Premier
Certified**

Tips and Tools for Network Availability and Compliance

Marty Adkins
Chesapeake NetCraftsmen, LLC

About the Speaker

- **Marty Adkins**
 - **Cisco CCIE #1289, CCSI #93021, Advanced Wireless and Data Center Application Services Design Specializations**
 - **Specialties: Large-Scale Routing & Switching, High Availability, Wireless LANs**
 - **Taught many of the Cisco courses plus some course development**
 - **Consultant to large federal and enterprise clients**

Agenda

- **Introduction**
- **Availability Drivers**
- **Software Components – Cisco IOS**
- **People, Processes and Tools**
- **Automation with Nectordia NetMRI**
- **Change and Outage Management**
- **Reducing MTTR**
- **Design and Management Tips**
- **Conclusion**

Today's Networks

- **Integral to success and performance of the organization – not just “plumbing”**
- **Many integrated services that are mission-critical**
 - Network & resource access control, encryption
 - Sophisticated DNS, DHCP
 - Highly accurate time service
 - Load balancing & caching, virtualization
 - WAN compression, acceleration
 - Unified communications, mobility, QoS

What Is “High Availability”?

- The **ability to define, achieve, and sustain** “target availability objectives” across services and/or technologies supported in the network **that align with the objectives of the business** (i.e. 99.9%, 99.99%, 99.999%)

Availability	Downtime per Year (24x7x365)		
99.000%	3 Days	15 Hours	36 Minutes
99.500%	1 Day	19 Hours	48 Minutes
99.900%		8 Hours	46 Minutes
99.950%		4 Hours	23 Minutes
99.990%			53 Minutes
99.999%			5 Minutes
99.9999%			30 Seconds

What Is “High Availability”?

- **Availability means more than just a total loss of service**
- **May be “in service” but not meeting formal SLA (or one inferred by the user)**
 - **Frame Relay CIR – resource is reachable but throughput / latency is unacceptable**
 - **Web server is reachable but is saturated due to load balancer error**
 - **Can place an IP voice call but it’s unintelligible**
- **User/customer definition: the proportion of time that a system can be used for productive work**

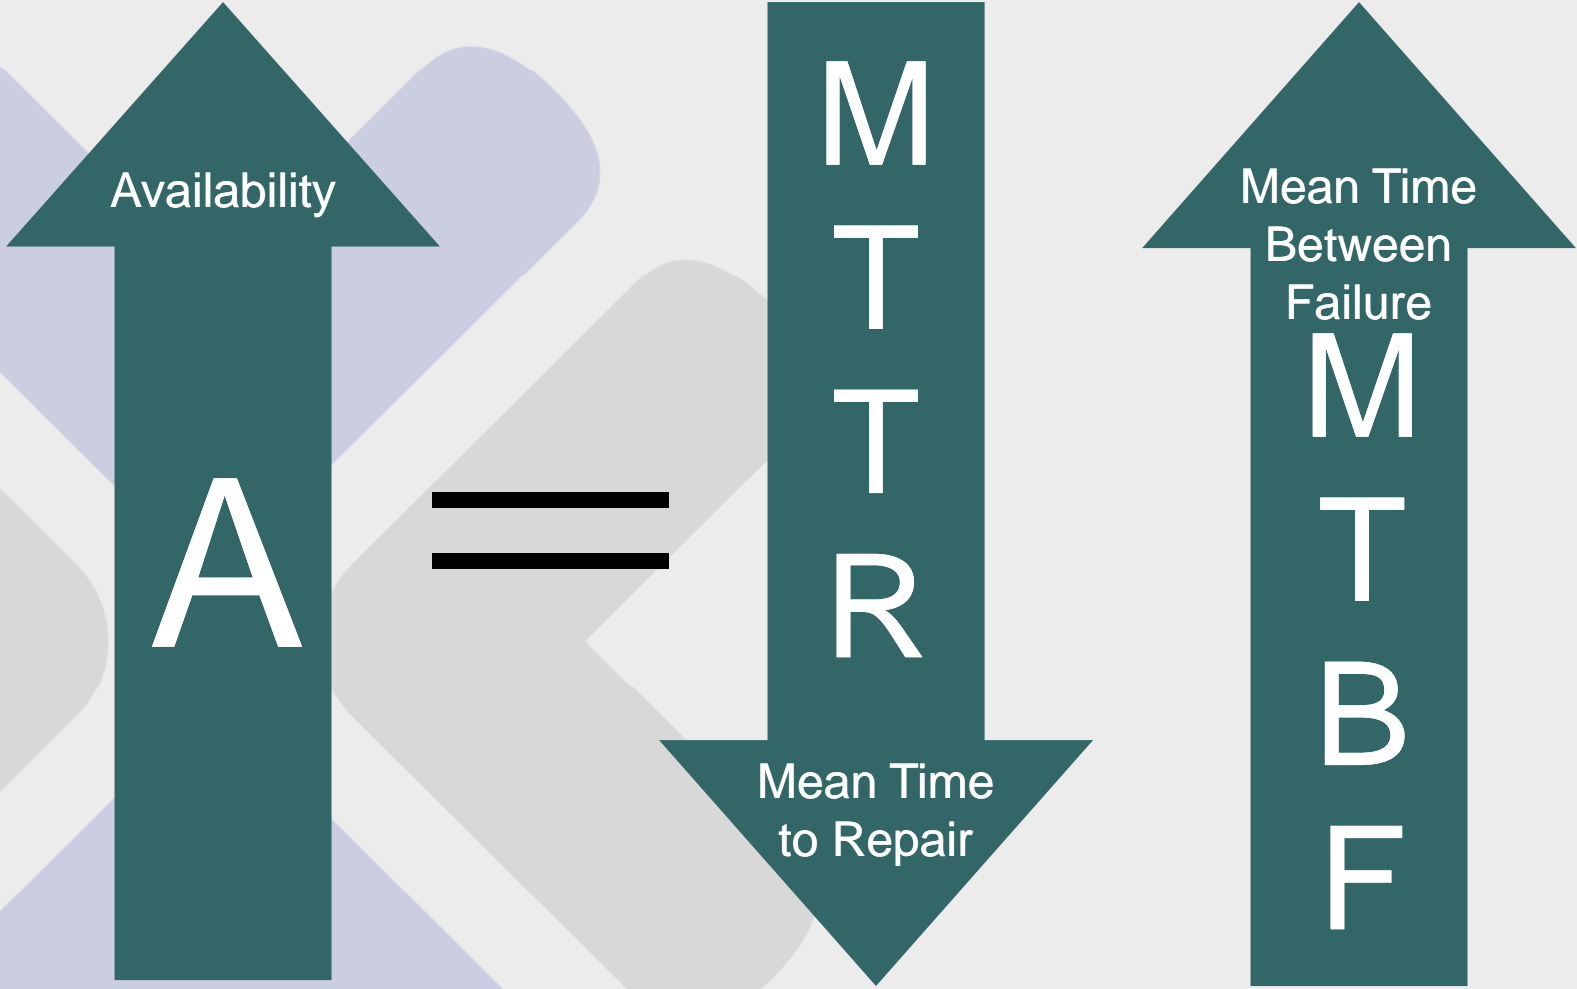
Classical Availability Calculations

- **Availability is calculated based on network design, component MTBF and MTTR**
- **MTBF = Mean Time Between Failure**
 - Calculated by measuring the average time between failures on a device or system
- **MTTR = Mean Time To Repair**
 - The time between when the device/network broke and when it was brought back into service
- **Availability = $MTBF / (MTBF + MTTR)$**

Availability Calculations – Revised

- **MTBF = Mean Time Between Failure**
 - Calculated by measuring the average time between failures on a device or system
- **MTTR = Mean Time To Repair**
 - The time between when the device/network broke and when it was brought back into service
- **MTTD = Mean Time To Detect**
 - The time between when the device/network broke and when it was noticed/detected (the *first* failure)
 - **This is a very real problem with redundant networks!**

Increasing Availability



Availability Demons

What Are the Time Bombs?

- **No technical ownership**
- **Layer 2/3 design**
- **Large failure domains**
- **Loose or non risk-aware change management**
- **High levels of network inconsistency**
- **Lack of network standards (SW, HW, config)**
- **No capacity planning or performance management**




Source: Cisco Systems 2004

Unscheduled Network Downtime Top Causes

- Change management
- Process consistency
- Methodology
- Communication

User Error
and Process
40%



Software and
Application
40%

Technology
20%

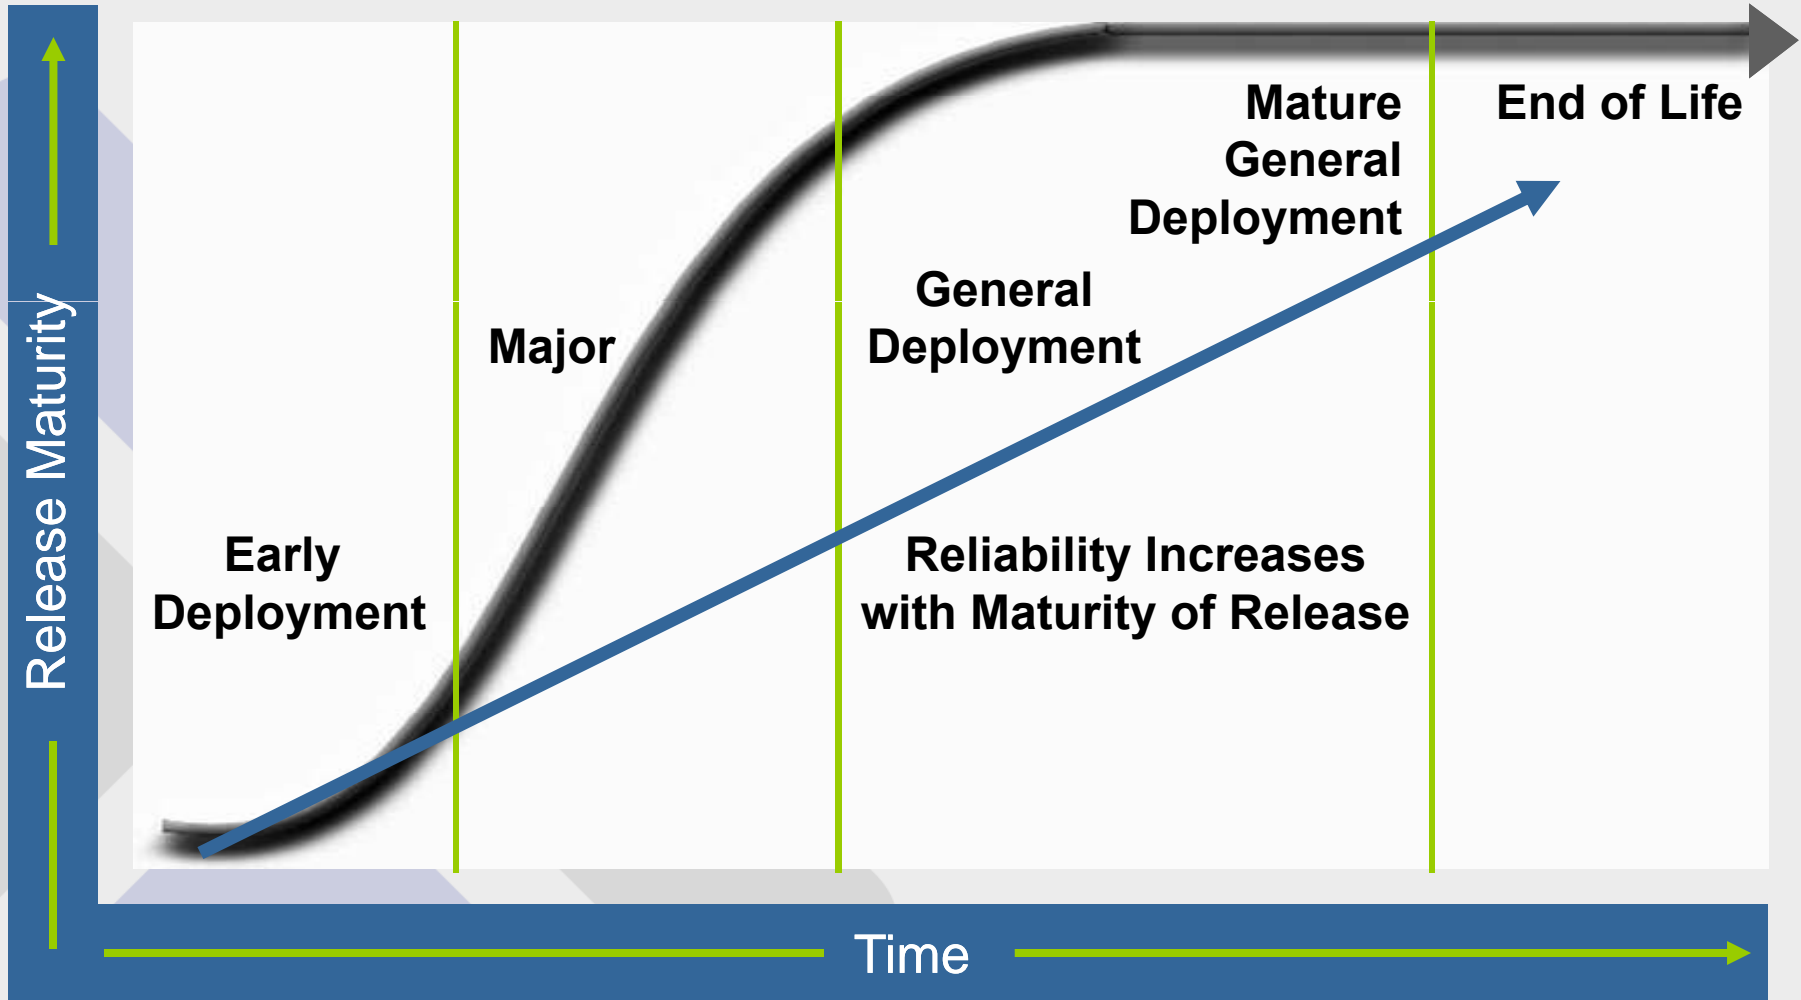
- Hardware
- Links
- Design
- Environmental issues
- Natural disasters

- Software issues
- Performance and load
- Scaling

Source: Gartner (2000)

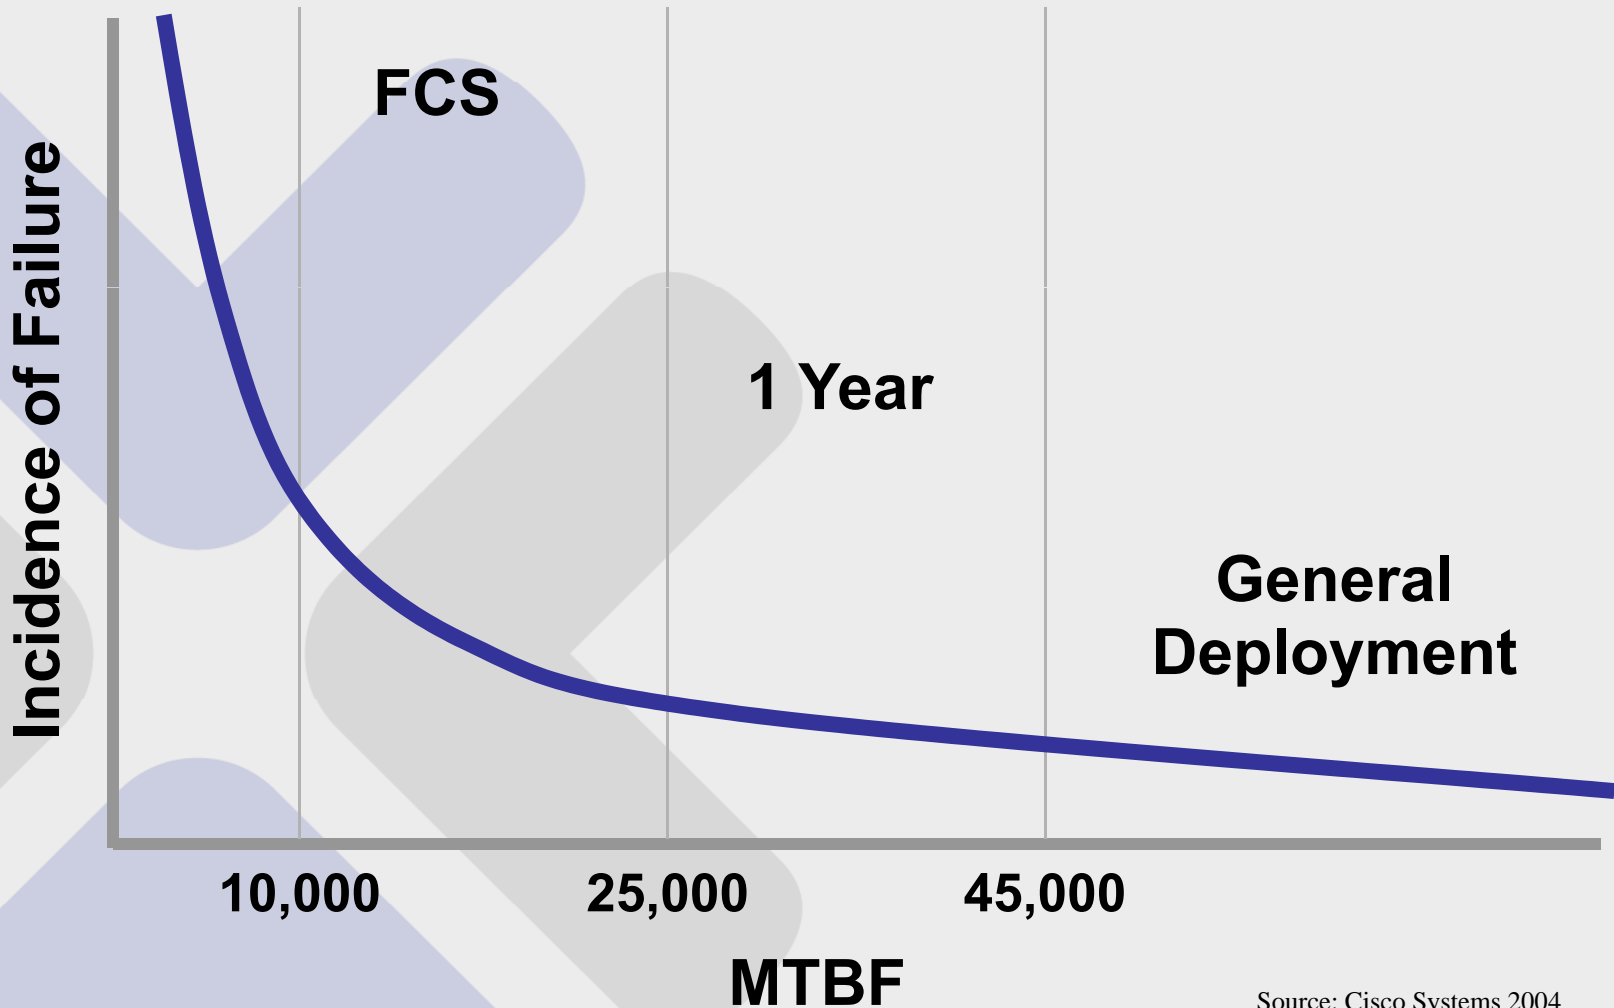
Software Reliability Factors

Age of Cisco IOS Release



Source: Cisco Systems 2004

Software Reliability Observed MTBF



Source: Cisco Systems 2004

Selecting an IOS Release Mainline

- **Mainline = stability (no new features), tends to become GD**
 - 12.4(25) 24th maintenance release of 12.4
 - 12.4(15b) First rebuild of 12.4(15)
 - 12.4(1) Initial release of 12.4
 - 12.4(8.3) Internal interim build →12.4(9)
 - 12.4(0.96) Internal beta build →12.4(1)
- **Rebuilds do not undergo full regression testing – limited to small patches**
- **So... which would you pick?**

Selecting an IOS Release 'T' Train

- **T-train = new features and platform support, begins as ED and progresses to LD**
 - **12.4(2)T4** Fourth rebuild of 12.4(2)T (new features)
 - **12.4(4)T5** Fifth rebuild of 12.4(4)T (more new features)
 - **12.4(11)T6** Sixth rebuild of 12.4(6)T (more new features)
 - **12.4(13)T3** Third rebuild of 12.4(13)T (more new features)
 - **12.4(15)T** First release of 12.4(15)T (more new features)
- **So... which would you pick?**


Selecting an IOS Release Evolution

- **12.3T matures and forms the basis for 12.4**
- **12.4T matures and forms the basis for... 15.0!**
- **So... would you pick 12.3T or 12.4?**
- **Cisco's recommendation**
 - **Try to get two solid releases for each train – current and previous**
 - **Only use T train where required for features or new platform support**
 - **Cisco Safe Harbor program for Cat 6500 family**
 - **Initially for financial services industry**
 - **Verifies IOS features, configs, topologies**
 - **Recently changed to engage at pre-FCS**

Unscheduled Network Downtime Top Causes

- Change management
- Process consistency
- Methodology
- Communication

**User Error
and Process
40%**



**Software and
Application
40%**

- Hardware
- Links
- Design
- Environmental issues
- Natural disasters


- Software issues
- Performance and load
- Scaling

Source: Gartner (2000)

What Keeps You Up at Night?

- **From where do you anticipate the greatest threat to network availability?**

“With 64% of our respondents’ votes, internal change continues to be network administrators’ greatest worry. While the vast majority of these changes are inadvertent, and probably trying to help the network, in the end they can end up hurting or crippling network performance.”



Netcordia survey of more than 450 network administrators,
December 2008

Internet Service Outages Top Causes

Oppenheimer, Ganapathi, Patterson of UC Berkeley in 2003:

"From a study of more than 500 component failures and dozens of user-visible failures in three large-scale Internet services, we observe that (1) **operator error** is the leading cause of failure in two of the three services studied, (2) **operator error** is the largest contributor to time to repair in two of the three services, (3) **configuration errors** are the largest category of operator errors..."

Internet Service Outages Top Causes

**A much less scientific polling on NANOG* in
2004:**

**Q: What configuration issues most affect the
performance and reliability of your network?**

A: Fingers... >;-)

***North American Network Operators' Group**

Overcoming The Availability “Wall”

- **Addressing 40% of network downtime – people, processes, and tools**
 - Hiring and training
 - IT process maturity
 - Automation
 - Change and problem management
- **Implement and *verify* best practices**
- **Prevent and/or rapidly detect “time bombs”?**
- **More sophisticated tools are required – we *must* move beyond ping and templates!**

Monitoring Network Change Compliance Assessment Methods

- **Automated tool grabs each device config on a schedule, or via asynchronous notification**
 - Syslog messages sent to collectors/monitors
 - Filtering on %CONFIG messages can trigger collection of the updated config plus *who* did it
 - Can maintain very detailed revision history
 - Requires login credentials
 - Tool compares text config snapshot to policy rules
- **Automated tool collects and analyzes operational data via SNMP (and possibly act on traps)**
- **Automated tool executes scripts that login to devices, inspect operation, perform heuristic analysis, generate reports and/or notification**

Netcordia NetMRI Capabilities

- **Assesses the impact of changes to the network for correctness and stability**
- **Automatically emulates what a team of network experts would do – diagnose, identify issues requiring review, and repair as directed**
- **Verifies compliance with industry/vendor best practices and enables customization as needed**
- **Policy scripts (when authorized) can modify active device configs to effect policy or remediate a problem**
- **Available as an appliance or as a VM**

NetMRI – How it Works


- **Expert analysis for ~275 heuristics based on SNMP-gathered data**
- **Tracks and reports on changes to network topology and operation**
- **Tracks and reports on changes to network device configurations**
 - **Maintains almost limitless configuration revisions and OS history**
- **Maintains 30 days of issue history and device data**
 - **post-mortem event correlation, trending**

Example – Civilian Agency


- **Large enterprise network:**
 - Campus HQ in Washington, D.C.
 - Hundreds of remote sites connected via VPN and Frame Relay
 - 2200 network devices, mostly Cisco
 - Supports >15,000 users
 - Stringent requirements for BCDR, COOP
- **Has utilized NetMRI Enterprise for 4+ years**

Entire Network 2009-02-03 / Daily


Issues Changes Policy Compliance Performance VoIP Events



Info Count
20




Warning Count
29



Error Count
19

Overall Score



■ Configurations

■ Routing

■ VLANs

■ Devices

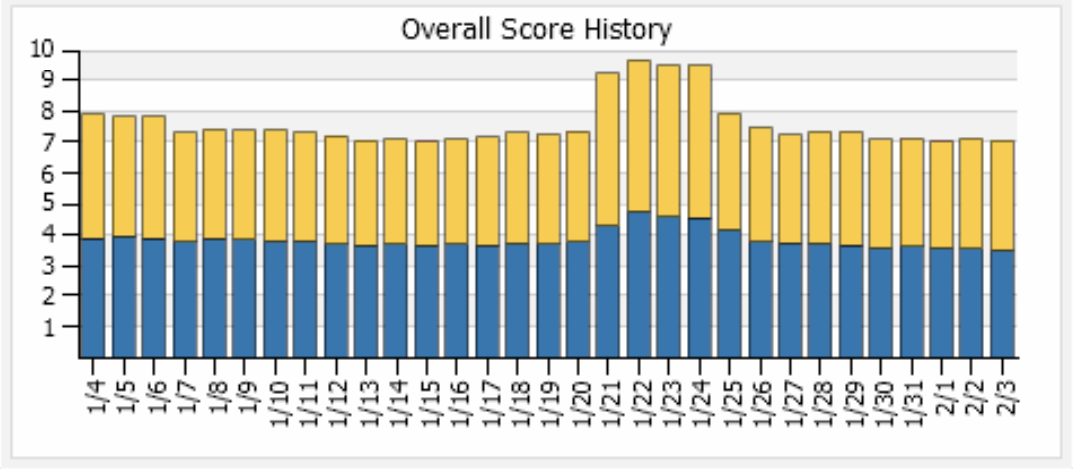
■ Security

■ VoIP

■ Interfaces

■ Subnets

■ Wireless



Displaying 1 - 68 of 68

Views Filters [Export] [Print] [Search]

Severity	Generated	Title	Component	# Affected	# New	# Condition
Error	2009-02-03 00:02:28.0	Trunk Port With PortFast Enabled	Interfaces	16	0	16
Error	2009-02-03 00:02:35.0	Access Port With PortFast Disabled	Interfaces	3	3	0
Error	2009-02-03 00:27:31.0	VLAN Trunk Port Down	VLANs	6	0	0
Error	2009-02-03 00:33:11.0	HSRP Not Recognizing Peer	Routing	12	0	0
Error	2009-02-03 00:37:59.0	VLAN Member Priority	VLANs	4	0	0

Page 1 of 1

Issues by Type Issues by Device

Example – Civilian Agency

- **NetMRI fired real-time issue of “VLAN Trunk Port Down” for a Cisco Catalyst 6500 switch...**
- **Access layer 2950 switch had dual fiber uplinks to two 6500 switches but...**



VLAN Trunk Port Down

Showing details for entire network

In:



Component:	VLANs	Correctness:	0
Severity:	Error	Analysis Start:	N/A (Realtime)
Generated:	2008-12-16 18:24:20	Analysis End:	N/A (Realtime)
Stability:	-2	Analysis Task:	RealTime Analysis Issue Definitions

Components Affected by Issue (Unsuppressed)

Displaying 1 - 1 of 1

Views Filters


<input type="checkbox"/>	IP Address	Device Name	Device Type	Interface	Timestamp	Diff	Sup?
<input type="checkbox"/>	10.2.23.2	aslsb3	Switch (99%)	3/13 - multi mode fiber fast ethernet	2008-12-16 18:17:34	Added	

Page 1 of 1

[Suppress Issues](#) | [Unsuppress Issues](#) | [Schedule Job](#) | [Execute Command](#)

History

Description





Type: Switch (99%) Vendor: Cisco
 O/S Version: 8.6(3) Model: wsc6506
 Up Time: 453d 09h 18m 07s SNMP Status: Enabled
 Last Update: 2009-02-03 11:36:22

Issues



2008/12/16

Period
Daily

Displaying 1 - 3 of 3

Views Filters

Severity	Generated	Title	Component	# Affected	# New	# Condition Change	# No Change	# Resolved	#
Error	2008-12-16 18:24:20.0	VLAN Trunk Port Down	VLANs	2	1	0	1	0	1
Warning	2008-12-16 00:14:09.0	CDP Neighbor Changed	Devices	33	8	25	0	2	0
Info	2008-12-16 23:52:38.0	Downstream Hub or Switch	Interfaces	441	5	0	436	2	0

Device

- Issues
- Changes
- Policy Compliance
- Custom Data
- Open Services
- Identification
- Performance
- CDP Neighbors
- Neighbors
- Location
- Inventory
- Environmental
- Logs
- Management Status
- Settings


- Config Files +
- Switch +
- Interfaces +

Example – Civilian Agency

- NetMRI fired real-time issue of “VLAN Trunk Port Down” for a Cisco Catalyst 6500 switch...
- Access layer switch had dual fiber uplinks to two 6500 switches but... one became inoperable
- **Cause: technician accidentally disturbed the wrong fiber patch. No service impact due to redundant uplinks**
- **NetMRI alert for time bomb – tech “defused”**
- **Follow up inspection showed port was up/up, running with no errors**

aslsb3 | 3/13 - multi mode fiber fast ethernet

ifIndex: 40 Device IP: 10.2.23.2
Type: ethernet-csmacd MAC Address: 00:02:FC:E1:8B:80
Speed: 100Mbps / fullDuplex / trunking Interface IP(s):
Status: up / up as of 2008-12-17 16:29:50.0 Port Fast: disabled



2008-12-16 / Daily

Measurement	Inbound			Outbound		
	Count	Rate	Percent	Count	Rate	Percent
Octets	1,794,769	166.18bps	2.0E-4	4,731,989	438.15bps	4.0E-4
Packets	12,854	0.1488/s	N/A	57,729	0.6682/s	N/A
Unicasts	7,473	0.0865/s	58.1375	4,384	0.0507/s	7.5941
Non-Unicasts	5,381	0.0623/s	41.8625	53,345	0.6174/s	92.4059
Multicasts	5,381	0.0623/s	41.8625	53,340	0.6174/s	92.3972
Broadcasts	0	0/s	0.0	5	0.0001/s	0.0087
Discards	2,791	0.0323/s	21.7131	0	0/s	0.0
Errors	0	0/s	0.0	0	0/s	0.0
Changes	0	0/s	0.0			
Alignment Errors	0	0/s	0.0			
FCS Errors	0	0/s	0.0			
Late Collisions				0	0/s	0

>>

Interface +

Performance -

- Summary
- Rates
- Percents
- Counts
- Charts

© 2009 Netcordia, Inc. All rights reserved.



Copyright 2010

Example – Civilian Agency

- **NetMRI fired real-time issue of “Switch Port Duplex Mismatch” for a Cisco 2950 switch port...**
- **Hint: Attached server had recently been promoted from test to production role by support vendor.**



Switch Port Duplex Mismatch

Showing details for entire network

In: Entire Network (2261)



Component: Interfaces **Correctness:** -2
Severity: Error **Analysis Start:** N/A (Realtime)
Generated: 2008-12-27 17:15:45 **Analysis End:** N/A (Realtime)
Stability: 0 **Analysis Task:** RealTime Analysis Issue Definitions

Components Affected by Issue (Unsuppressed)

Displaying 1 - 1 of 1


Views Filters

Device	Interface	Duplex Setting	Total Packets	% Errors	Neighbor	Timestamp	Diff	Sup?
10.4.26.57 2950crwv1	Fa0/4 - BigFix Enterprise Patcher	In Out	1,628,152 1,311,293	17.11455 0.00000	unknown	2008-12-27 17:14:06	Added	

Page 1 of 1

Suppress Issues | Unsuppress Issues | Schedule Job | Execute Command


History



2950crwv1 | Fa0/4 - BigFix Enterprise Patcher



ifIndex: 4 **Device IP:** 10.4.26.57
Type: ethernet-csmacd **MAC Address:** 00:0D:65:F5:52:C4
Speed: 100Mbps **Interface IP(s):**
Status: up / up as of 2009-01-28 16:32:27.0 **Port Fast:** enabled



2008-12-27 / Daily ?

Measurement	Inbound			Outbound		
	Count	Rate	Percent	Count	Rate	Percent
Octets	106,762,072,171	9.89Mbps	9.8854	5,737,609,625	531.26Kbps	0.5313
Packets	92,654,284	1072.3751/s	N/A	73,910,278	855.4331/s	N/A
Unicasts	92,653,921	1072.3709/s	99.9996	73,810,198	854.2748/s	99.8646
Non-Unicasts	363	0.0042/s	4.0E-4	100,080	1.1583/s	0.1354
Multicasts	0	0/s	0.0	98,267	1.1373/s	0.133
Broadcasts	363	0.0042/s	4.0E-4	1,814	0.021/s	0.0025
Discards	0	0/s	0.0	0	0/s	0.0
Errors	20,767,122	240.3574/s	18.3097	0	0/s	0.0
Changes	0	0/s	0.0			
Alignment Errors	0	0/s	0.0			
FCS Errors	20,748,150	240.1378/s	99.9086			
Late Collisions				0	0/s	0

>>

Interface +

Performance -

- ☰ Summary
- ☰ Rates
- ☰ Percents
- ☰ Counts
- ☰ Charts

© 2009 Netcordia, Inc. All rights reserved.



Copyright 2010

Example – Civilian Agency

- NetMRI fired real-time issue of “Switch Port Duplex Mismatch” for a Cisco 2950 switch port...
- Hint: Attached server had recently been promoted from test to production role by support vendor.
- **Solution: Vendor reinitialized NIC settings to auto/auto; also verified NIC teaming configuration at network staff’s request**

Example – Civilian Agency

- **NetMRI cited remote office router for “Config Running Not Saved” issue**
- **Issue details include timestamps of last reboot, last configuration change, last save**
- **Network administrator confirmed that the changes were required and needed to be made permanent**
- **An approved NetMRI script was invoked to perform that according to the site policy (including back up to TFTP server)**

Config Running Not Saved

Showing details for entire network

In: Entire Network (2265)



Component: Configurations Correctness: -0.5
Severity: Info Analysis Start: N/A (Realtime)
Generated: 2009-02-02 12:43:31 Analysis End: N/A (Realtime)
Stability: 0 Analysis Task: RealTime Analysis Issue Definitions

Components Affected by Issue (Unsuppressed)

Displaying 1 - 1 of 1


Views Filters

IP Address	Device Name	Reboot Time	Changed Time	Saved Time	Time Difference	Timestamp	Diff	Sup?
10.40.212.2	albuquerque	2009-01-05 15:14:53	2009-01-08 10:42:45	2009-01-05 15:14:53	2d 19:27:52	2009-02-02 00:08:40	Same	

Page 1 of 1

Suppress Issues | Unsuppress Issues | Schedule Job | Execute Command

History





Config Running Not Saved

Showing details for entire network

In: Entire Network (2265)



Component: Configurations **Correctness:** -0.5
Severity: Info **Analysis Start:** N/A (Realtime)
Generated: 2009-02-02 12:43:31 **Analysis End:** N/A (Realtime)
Stability: 0 **Analysis Task:** RealTime Analysis Issue Definitions

Components Affected by Issue (Unsuppressed)

Displaying 1 - 1 of 1



Views Filters

<input type="checkbox"/>	IP Address	Device Name	Reboot Time	Changed Time	Saved Time	Time Difference	Timestamp	Diff	Sup?
<input checked="" type="checkbox"/>	10.40.212.2	albuquerque	2009-01-05 15:14:53	2009-01-08 10:42:45	2009-01-05 15:14:53	2d 19:27:52	2009-02-02 00:08:40	Same	

Page 1 of 1

Suppress Issues | Unsuppress Issues | Schedule Job | Execute Command

History



Edit Job

This form is used to create a job specification that details a specific script to be run on a certain schedule against selected devices.

Job Name: Approved:

Script Name:

Description:

Once on February 02 at 16:40

Device Groups

Devices

..... albuquerque (10.40.212.2)

entire Network (2265)

?

Views Filters

Time Difference	Timestamp	Diff	Sup?
2d 19:27:52	2009-02-02 00:08:40	Same	

Example – Civilian Agency

- **When real-time notification really matters...**
- **All diagnostic issues may optionally trigger real-time notification – email, Syslog, SNMP trap**
- **HTML-formatted mail message includes details plus link to issue details**

NetMRI Issue Notification

Network Name: Govnet

Server Name: netmri

Serial No: 2712-50001

Run Date: 2009-01-17 05:21:09

Start Date: 2009-01-17 00:50:59

End Date: 2009-01-17 05:21:09

Severity	Issue	Instances	Timestamp
Error	Device Power Supply Failure	1	2009-01-17 05:20:24

DevicePowerError

IP Address	Device Name	Description	State	Low Shutdown	High Shutdown	Timestamp	DiffState
10.6.112.50	asWAS-3a	Power Supply 1, WS-CAC-2500W	critical	N/A	N/A	2009-01-17 05:18:59	added

NetMRI Issue Notification

Network Name: Govnet

Server Name: netmri

Serial No: 2712-50001

Run Date: 2009-01-11 00:00:49

Start Date: 2009-01-10 00:27:34

End Date: 2009-01-11 00:00:49

Severity	Issue	Instances	Timestamp
Warning	Device Fan Warning	1	2009-01-10 23:59:57

DeviceFanWarning

IP Address	Device Name	Description	State	Low Warning	High Warning	Timestamp	DiffState
10.3.97.22	2950c-BldgC-2	chassis	warning	N/A	N/A	2009-01-10 23:48:44	added

Example – Civilian Agency

- **Support for policy compliance ...**
- **Regularly scheduled Inspector General audit:**
 - **Q: Can you track device configuration changes going back 12 months?**
 - **A: Sure**
 - **IG auditor selected a device**
 - **NetMRI had configuration versions back to mid-2006, when the device was first installed.**
 - **IG was provided a PDF showing the side-by-side differences due to an authorized change**

Comparing Configuration Files

Selected Files

Host Device: [10.1.31.1 \(arc431\)](#)



File Status: Archived Running

Last Modified: 2008-11-18 13:07:54 by Unknown

Host Device: [10.1.31.1 \(arc431\)](#)



File Status: Archived Running

Last Modified: 2008-11-18 15:53:40 by Unknown

Find:

view

action

Changes: 0

Removals: 0

Additions: 3

```
1 version 12.2
2 no service pad
3 service tcp-keepalives-in
4 service tcp-keepalives-out
5 service timestamps debug datetime localtime show-timezone
6 Skipping to line 56
7 ip address 10.1.39.3 255.255.255.0
8 no ip proxy-arp
9 ip pim sparse-dense-mode
10 ip multicast ttl-threshold 255
11 ip multicast boundary No_Local_Scope
12
13 !
14 interface Vlan55
15 description arC431-drCSW2
16 ip address 10.1.55.3 255.255.255.0
17 no ip proxy-arp
18 ip pim sparse-dense-mode
19 ip multicast ttl-threshold 255
20 ip multicast boundary No_Local_Scope
21
22 !
23 interface Vlan303
```

```
1 version 12.2
2 no service pad
3 service tcp-keepalives-in
4 service tcp-keepalives-out
5 service timestamps debug datetime localtime show-timezone
6 Skipping to line 56
7 ip address 10.1.39.3 255.255.255.0
8 no ip proxy-arp
9 ip pim sparse-dense-mode
10 ip multicast ttl-threshold 255
11 ip multicast boundary No_Local_Scope
12 arp timeout 300
13
14 !
15 interface Vlan55
16 description arC431-drCSW2
17 ip address 10.1.55.3 255.255.255.0
18 no ip proxy-arp
19 ip pim sparse-dense-mode
20 ip multicast ttl-threshold 255
21 ip multicast boundary No_Local_Scope
22 arp timeout 300
23
24 !
25 interface Vlan303
```

Comparing Configuration Files

Selected Files

Host Device: [10.1.31.1 \(arc431\)](#)
File Status: Archived Running
Last Modified: 2008-11-18 13:07:54 by Unknown

Host Device: [10.1.31.1 \(arc431\)](#)
File Status: Archived Running
Last Modified: 2008-11-18 15:53:40 by Unknown

Find:

view

action

Changes: 0

Removals: 0

Additions: 3

- Side-by-Side
- Inline
- Over/Under
- Entire File
- Changes Only
- Number of lines to show around changes ▶
- Swap Files

```
1 version 12.2
2 no service pad
3 service tcp-keepalives-in
4 service tcp-keepalives-out
5 service timestamps debug d
Skipping to line 56
56 ip address 10.1.39.3 255.255.255.0
57 no ip proxy-arp
58 ip pim sparse-dense-mode
59 ip multicast ttl-threshol
60 ip multicast boundary No
61
62 !
63 interface Vlan55
64 description arc431-drCSW2
65 ip address 10.1.55.3 255.255.255.0
66 no ip proxy-arp
67 ip pim sparse-dense-mode
68 ip multicast ttl-threshold 255
69 ip multicast boundary No_Local_Scope
70
71 !
72 interface Vlan303
```

```
1 version 12.2
2 no service pad
3 service tcp-keepalives-in
4 service tcp-keepalives-out
5 service timestamps debug datetime localtime show-timezone
Skipping to line 56
56 ip address 10.1.39.3 255.255.255.0
57 no ip proxy-arp
58 ip pim sparse-dense-mode
59 ip multicast ttl-threshold 255
60 ip multicast boundary No_Local_Scope
61 arp timeout 300
62 !
63 interface Vlan55
64 description arc431-drCSW2
65 ip address 10.1.55.3 255.255.255.0
66 no ip proxy-arp
67 ip pim sparse-dense-mode
68 ip multicast ttl-threshold 255
69 ip multicast boundary No_Local_Scope
70 arp timeout 300
71 !
72 interface Vlan303
```

Example – Civilian Agency

- **Cisco IOS device code push – making code updates bullet proof ...**
- **Upgraded IOS in 800+ switches via NetMRI script:**
 - Checked if new image file already present
 - Checked flash memory for available space – if necessary, deleted non-active image file
 - Copied new image file to flash and compared MD5 hash value to published value on Cisco CCO
 - Changed config boot string to new image file
 - Generated custom issue report with detailed job status for each device

Run Configuration Command Script

The command will be run again these devices:
2960tf41n2

The selected configuration command script defines Script-Variables that require user input. Please provide the values for the fields specified below and then press the OK button to start the batch.

newimagename:

newimagemd5:

tftpserver:

OK

Reset

Cancel

Job Viewer

BatchID: 1934 Start Time: 2009-02-04 18:20:06
Script: 2960IOSUpgrade.ccs End Time: 2009-02-04 18:22:20
Job Count: 1 Status: OK



Details

Issues

Files

Job Details 2009/02/04

Refresh
Off

Ad Hoc Job 02/04 18:20 - 2960IOSUpgrade.ccs

Displaying 1 - 1 of 1

Views Filters

	Status	Start Time	End Time	IP Address	Device Name	Actions
	OK	2009-02-04 18:20:06	2009-02-04 18:22:14	10.3.24.85	2960tf41n2	

Page 1 of 1

Cancel All

Rerun Errors

Reschedule Errors

Job Viewer

BatchID: 1934 Start Time: 2009-02-04 18:20:06
Script: 2960IOSUpgrade.ccs End Time: 2009-02-04 18:22:20
Job Count: 1 Status: OK



Details

Issues

Files

Job Issues

Displaying 1 - 1 of 1

Views Filters

	Severity	Generated	Title	Component	# Affected	# New	# Condition Change
	Info	2009-02-04 18:22:15.0	c2960 IOS Upgrade Completed	Configurations	1	0	0

A new IOS image has been downloaded into the following c2960 switches. The listing shows the image file that was copied (if not already present), the new value of the 'boot system' string, and any previous image file that was deleted to make room.



c2960 IOS Upgrade Completed

Showing details for entire network

In: Entire Network (2264)



Component:	Configurations	Correctness:	0
Severity:	Info	Analysis Start:	2009-02-04 18:05:17
Generated:	2009-02-04 18:22:15	Analysis End:	2009-02-04 18:22:14
Stability:	0	Analysis Task:	Configuration Command Scripts

Components Affected by Issue (Unsuppressed)

Displaying 1 - 1 of 1

Views Filters

<input type="checkbox"/>	DeletedImage	BootImageName	ImageCopied	Host	Device Name
<input type="checkbox"/>	none	c2960-lanbasek9-mz.122-35.SE5.bin	c2960-lanbasek9-mz.122-35.SE5.bin	10.3.24.85	2960tf41n2




Page 1 of 1

Suppress Issues | Unsuppress Issues | Schedule Job | Execute Command

History

Description



Change Management

- **CM doesn't have to be totally bureaucratic**
 - Proposed changes can be circulated via email to other support team leads and management
 - Complicated changes that require coordinated changes on servers or other devices, merit meetings to flag potential misunderstandings
 - Once team leads agree, notice is given to end users
- ***Always* have a back out plan**
- **Post a follow up status after work is complete**
- **The staff that performed the work should still be on site for several hours after users resume work**

Outage Management Communication Is Key

- **As soon as the impact and scope is understood, notify *only IT teams* via email:**
 - “An Internet slow-down has been detected and NetOps is investigating.”
- **Important for others to not chase their tails, cause confusion, or give incorrect or contradictory explanations to users who may already be inquiring.**

Outage Management Informing Users/Customers

- **If the outage persists longer than X minutes, inform the user community succinctly:**
 - “We are experiencing an Internet slow-down. NetOps is working with the Internet service provider to determine the cause.”
- **Prevents additional needless user reports**
- **Helps network staff stay focused on a resolution**
- **Establishes ownership – we're on it!**

Outage Management Status Updates

- **If it's an extended outage, send a status update every Y hours:**
 - “NetOps continues to work with RIM engineers to solve the Blackberry outage. No estimate to repair is yet known.”
- **When solved, post an “all clear” message:**
 - “Blackberry service has been fully restored. If you are still experiencing an issue, please contact the IT Help Desk for assistance.”
- **Users perceive proactive accurate notification**
 - “No need to keep calling or emailing – they'll tell us when something changes or more is known.”

Outage Management Notification Methods

- **Email (to mobile devices) but beware of catch-22**
- **SMS text messaging via SNPP IP connection, cellular data or dial-out modem**
- **Twitter**
- **Voice call – manual or automated**

Decreasing MTTR

- **Staffing - at least two persons knowledgeable in each area; coordinate absences**
- **Redundant VPN and/or dial up access**
 - Separate VPN group address pool with access to network devices
 - A hardened bastion host requiring SSH (with token)
 - Remote desktop (but protect it!)
- **Copies of documentation and configs in electronic form, possibly at home**
- **Staff training !!! Good processes are necessary but insufficient**

Reducing MTTR

Out-Of-Band Management

- **Connect consoles of all devices to commservers (26xx with async ports)**
- **Use banners or menus to identify connected devices**
- **Commserver must authenticate users because the attached console device might not**
- **Avoid catch-22 issues – can't reach commserver due to network outage**
 - **Carefully factor in network redundancy and points of attachment**
 - **Consider building a separate mostly 'flat' network**
 - **Don't depend on DNS working – either connect by IP or have a separate DNS server for this purpose**

Decreasing MTTR

Value Of A Test Lab

- **All spare modules are burned in, running production code, ready to go**
 - RMA can take 24-48 hours; what if it's DOA?
- **Test new code and/or features in representative topology and configurations**
- **“Torture track” testing**
 - Flapping and one-way links
 - Marginal cables (optical attenuators)
- **Ideal training environment for staff**
- **Use commserver for ease of console access**

IP Address Planning

- **Must be hierarchical to support summarization**
- **Keep network device (plumbing) addresses totally separate from user space**
 - Use private IP space, e.g. 10.0.0.0/8
 - Greatly simplifies ACLs for SSH, SNMP, server wrappers
 - Helps limit address ranges for NMS to scan/discover
 - Requires a separate management VLAN on switches and trunking to edge switches
- **A single address allocation mechanism is a must (even if it's just a protected spreadsheet)**
 - Forms the basis for hierarchical allocation and delegation
 - Identifies the location/purpose of each CIDR block
 - Used by interdisciplinary teams to troubleshoot so store it in a shared network location

Device Naming Conventions

- **Names should include**

- Geographical location (requires a building and city scheme)
- Role in the network hierarchy

Examples:

BWI_coreA = the first core router in Baltimore

drC121 = 1st distribution router in bldg C, 1st floor, riser 2

sjc12-31-sw2 = San Jose bldg 12, 3rd floor, 1st IDF, switch 2

- **Names should almost never include model numbers**

- Will the names in a traceroute display be helpful?
- If you upgrade the device to a different model, will you change its name (and in DNS, TACACS/RADIUS, NMS)?
- How will you reconcile the discontinuity in syslog, baselining data, any databases containing the name, etc.?

Device Naming Conventions (cont.)

- **Router names in DNS**

- Canonical name is mapped to the loopback address (or management VLAN address)
- Each interface or VLAN name is formed by concatenation

Examples:

BWI_coreA-ge0-0 (GigabitEthernet0/0)

drC121-v121 (VLAN 121)

sjc12-31-sw2 (it's not a router so it only has one IP and name!)

- **Traceroutes will be lucid, including for load-shared paths**
- **Remember to source all management traffic from the loopback or management VLAN address**
 - Syslog, SNMP traps, TFTP, TACACS/RADIUS, NTP

What To Include In Network Documentation

- **Logical and physical network drawings**
 - Layer 1 – physical cable tracing, fiber pairs, patches, etc.
 - Layer 2 - VLANs, STP roots, blocking ports, etherchannels, VACLs, QoS
 - Layer 3 – Addressing, HSRP roles, loopbacks, null0 routes, default routes, summarization, redistribution, route filtering, tweaking of metrics, vrf, QoS, ACL traffic filters
- **Definitely requires multiple drawings, following the hierarchy and subsystems**
- **Mandate updated designs & drawings prior to approval of a change or new installation. Much better to catch errors up front**

What To Include In Network Documentation (cont.)

- **ACLs and/or firewall rules**
 - Annotate in a tracking document
 - Explain purpose for each statement; identify IP addresses
 - Enter date an entry was added, by whom, and the requester and approver.
 - Move deleted entries to a separate section at the end
- **ACL configurations**
 - Use named ACLs and include the date – Inet_In_013110
 - Activate by *replacing* the “ip access-group” statement
 - Monitor ACE match counts and/or logs
 - Keep one previous ACL version in config so you can quickly revert back.

Managing Device Configurations

- **Create annotated config templates which explain rationale for statements and values**
 - Useful as a baseline for verification/audit
 - Perform knowledge transfer, including to new staff or vendors
 - Helps *you* remember why you did something a year from now
- **Auditing live configs vs. best practice templates is a hard problem**
 - Too tedious and error prone, so it doesn't get done
 - Cisco NCM and Netcordia NetMRI perform OOTB and custom policy-based audits via wizard interface
 - Consider open source tools – you tailor and support
 - Poor man's method – grep, diff

Backing Up Device Configurations

- **After completing config changes, immediately back up configs to TFTP/SFTP server**
 - Need consistent naming scheme or else use Cisco autogenerated names
 - TFTP server should have wrappers to limit access to only network device and network staff IPs
 - Cisco NCM and Netcordia NetMRI will grab config upon receipt of Syslog “%CONFIG” message
 - Back up configs to removable flash in case device needs to be swapped out – lowers MTTR

Syslog Tips

- Use more than one Syslog server, physically diverse – consider syslog-ng and Splunk
- Use wrappers to control access by source IP
- Login only via SSH2 and/or token instead of Telnet
- To scale, filter the Syslog traps on the fly into the appropriate file, based on facility code or device partial name (a good naming convention helps!)
- Have a housekeeping job back up the log files; and perform a version roll over once a week
 - ciscolog.1 = current
 - ciscolog.2 = previous week, etc.
- Use grep or a script to search logs for strings of interest (severity, facility code, wildcarded name)

Does Anybody Really Know What Time It Is?

- **Use of NTP in all devices is critical to accurate event correlation, also for Windows AD / Kerberos**
- **Build a distribution tree**
 - Buy a time receiver (GPS or CDMA)
 - Have two or more routers poll public NTP servers
 - Other routers poll them
 - Use “ntp server x.x.x.x prefer” to prevent client oscillation
 - Switches (and user hosts) poll their default gateway
- **No Internet access – use IOS as an NTP master**
 - Accuracy isn't as important as precision (being in lock step)
- **Use NTP authentication for network devices**
- **Configure IOS timestamps for 'datetime'**

Best Practices Security Related

- **Implement Reverse Path Forwarding (RPF) check on all edge LANs**
 - It's really difficult to track down an infected PC with a source of 169.254.100.77
 - Shield the rest of the world from the crud
 - Use RPF ACL to create exception for asymmetric HSRP
- **Use DHCP snooping to prevent rogue DHCP servers**
- **On high traffic routers, instead of ACL logging, use netflow and “ip accounting access-violations”**
- **Use AAA to log privileged commands to ACS server – useful for post mortem (and staff becomes more careful)**
- **Remember Vista / Win7 and MacOS try IPv6 first**

Take Away Thoughts...

- **Hardware and software continue to improve for both MTBF and MTTR. The real drivers mostly depend on us.**
- **A scalable modular repeatable (boring) design goes a long way towards improving availability.**
- **Redundancy is good. Too much redundancy stresses protocols in areas where they previously have not received much exercise.**
- **Redundancy has not only a hardware cost, but a human cost, and one must commit to train the humans to understand it well, lest they greatly impact availability.**

Take Away Thoughts...

- **Proactive network assessment is hard – it requires sophisticated tools that better approximate the seasoned eye of a network expert with infinite time.**
- **If you don't measure and you don't have a baseline, then you have no idea where you've been, and no way to judge new territory.**
- **You must track and control change!**


Take Away Thoughts...

I spent a few hours talking with some leading network managers and was surprised many of them have the same approach when it comes to issues or outages. Most of them seem to focus more time and resources on trying to shorten the time to troubleshoot and resolve instead of trying to eliminate the problem from occurring in the first place. One favorite quote I heard was “I can prove if we shorten the MTTR from 3 hours to 90 minutes and my boss loves that. But it's almost impossible for me to prove we avoided X number of major outages by proactively managing network maintenance including change and configuration. We all know it makes sense, but I can't quantify it, so it gets put on the back burner.”

- Matt Gowarty of Netcordia

Any Questions?



- For a presentation copy, please email madkins@netcraftsmen.net
- About Chesapeake NetCraftsmen:
 - Cisco Premier Partner
 - Cisco Customer Satisfaction **Excellence** rating 
 - Developed numerous courses for Cisco (internal and public)
- Cisco Advanced Specializations
 - Advanced Routing and Switching (12+ CCIEs on staff)
 - Advanced Security (four double R&S/Security CCIEs so far)
 - Advanced Unified Communications (and IP Telephony)
 - Advanced Wireless
 - Advanced Data Center
- Knowledge transfer is key to every project!



CISCO™

PARTNER

**Premier
Certified**

Extra Slides

One Highly-Available Device vs. Two Simpler Ones

- Life cycle cost is higher for two devices
- IOS fails more often than hardware
- Ease of getting maintenance window
- Code upgrades – can do one chassis at a time
- Failover times
 - Intrachassis – IOS RPR/SSO
 - Interchassis – L2 & L3 protocols, timers
- Either way, staff must understand how it works!

Decreasing MTTR

Tips For The Finger Tips

- The IOS command “|” pipe operator is a great time saver – learn to use it

```
RouterA>sh int | i line|minute
Ethernet0 is up, line protocol is up
  5 minute input rate 0 bits/sec, 1 packets/sec
  5 minute output rate 1000 bits/sec, 1 packets/sec
Ethernet1 is up, line protocol is up
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
Serial0 is administratively down, line protocol is down
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
Serial1 is administratively down, line protocol is down
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
```

Decreasing MTTR

Tips For The Finger Tips

- The IOS command “|” pipe operator is a great time saver – learn to use it

```
RouterA>show proc cpu | e 0.00
```

```
CPU utilization for five seconds: 3%/0%; one minute: 1%; five minutes: 1%
```

PID	Runtime(ms)	Invoked	uSecs	5Sec	1Min	5Min	TTY	Process
3	5752	262	21954	1.47%	0.13%	0.44%	2	Virtual Exec
22	13519864	2567777	5265	1.22%	0.12%	0.06%	0	IP Input

```
2950CC3161>sh int status | i half
```

```
Fa0/5          connected    340          a-half    a-100 10/100BaseTX
```

Decreasing MTTR

Tips For The Finger Tips

- The IOS command “|” pipe operator is a great time saver – learn to use it

You need to change a config statement on 200 subinterfaces

```
FrameHub1#sh run | i terface\ Hssi2/0\.|verify
interface Hssi2/0.101 point-to-point
ip verify unicast reverse-path 121
interface Hssi2/0.102 point-to-point
ip verify unicast reverse-path 121
interface Hssi2/0.103 point-to-point
ip verify unicast reverse-path 121
```

**Copy/paste into your favorite editor. Do a global search/replace.
Copy/paste the result back into the router config.**