Netcordia, Inc.

Silly Network Management Tricks

Terry Slattery
Founder

Netcordia, Inc.

Network Management Silliness Abounds

- Based on ...
  - Consulting experience
  - Our experience in developing **NetMRI**
- Modeled after Bill Cheswick's "Silly Networking Tricks" presentation

Network Management Perspective

- 1990 - HP OpenView
- Plethora of products since then
  - Seldom providing what network people need
  - Many, many point products
- Are you happy with current network management?

2005 - Not much progress

- HPOV still around
- CiscoWorks
- Abundance of point products
  - Wireless; VoIP; Route analysis
  - Performance reporting
  - Fault reporting
- Multiple tools per shop
  - HPOV, CiscoWorks suite
  - Concord, Airwave, NetIQ
- Learning all the systems?
- Questionable usefulness together

Silliness Topics

- User interfaces
- Basic premises
- Data vs information
- Polling
- Network Configuration

UI - Event Managers

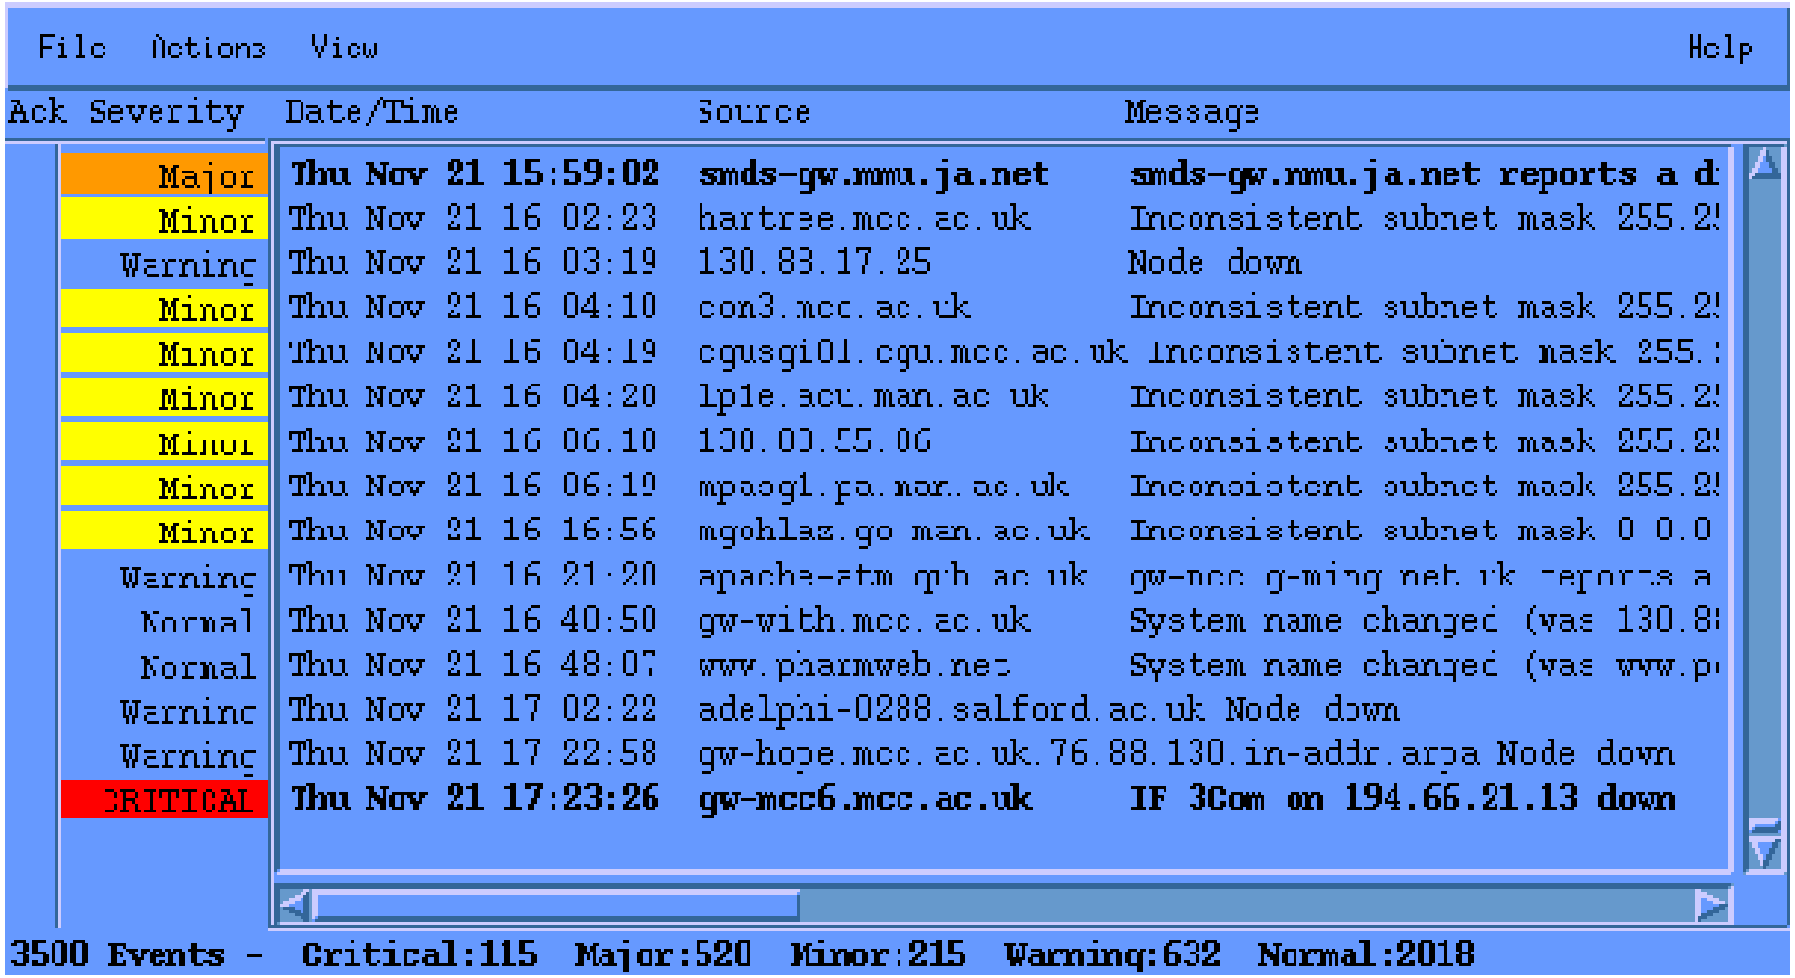

UI - Event Managers

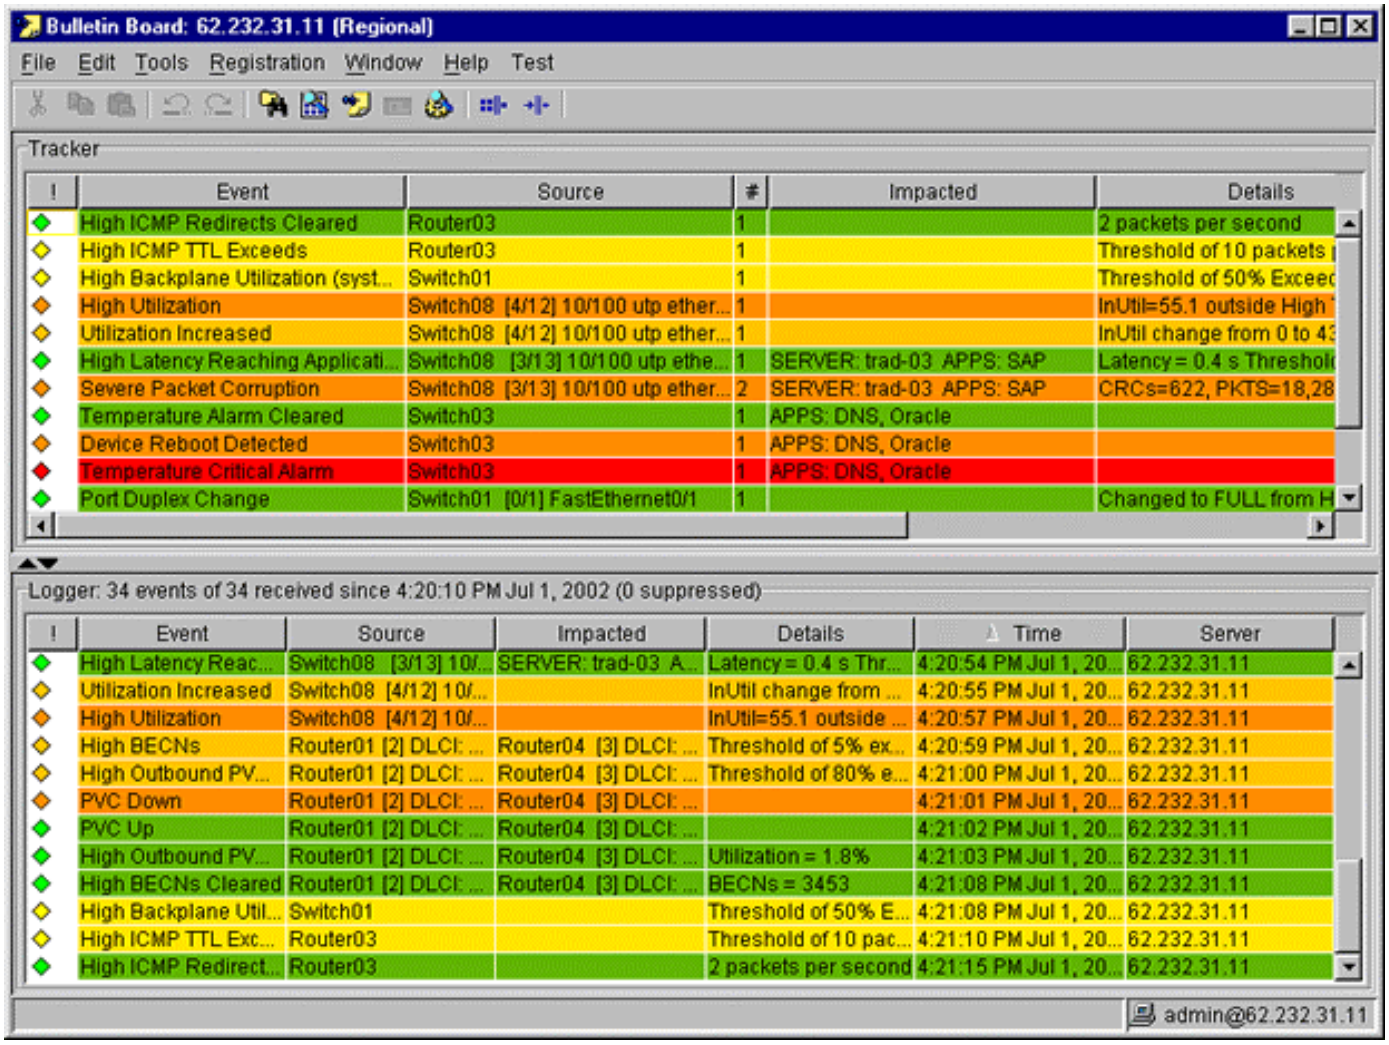

UI - Event Management

- Error classification is more important than timestamp
- Group by severity
- Sub-group by error

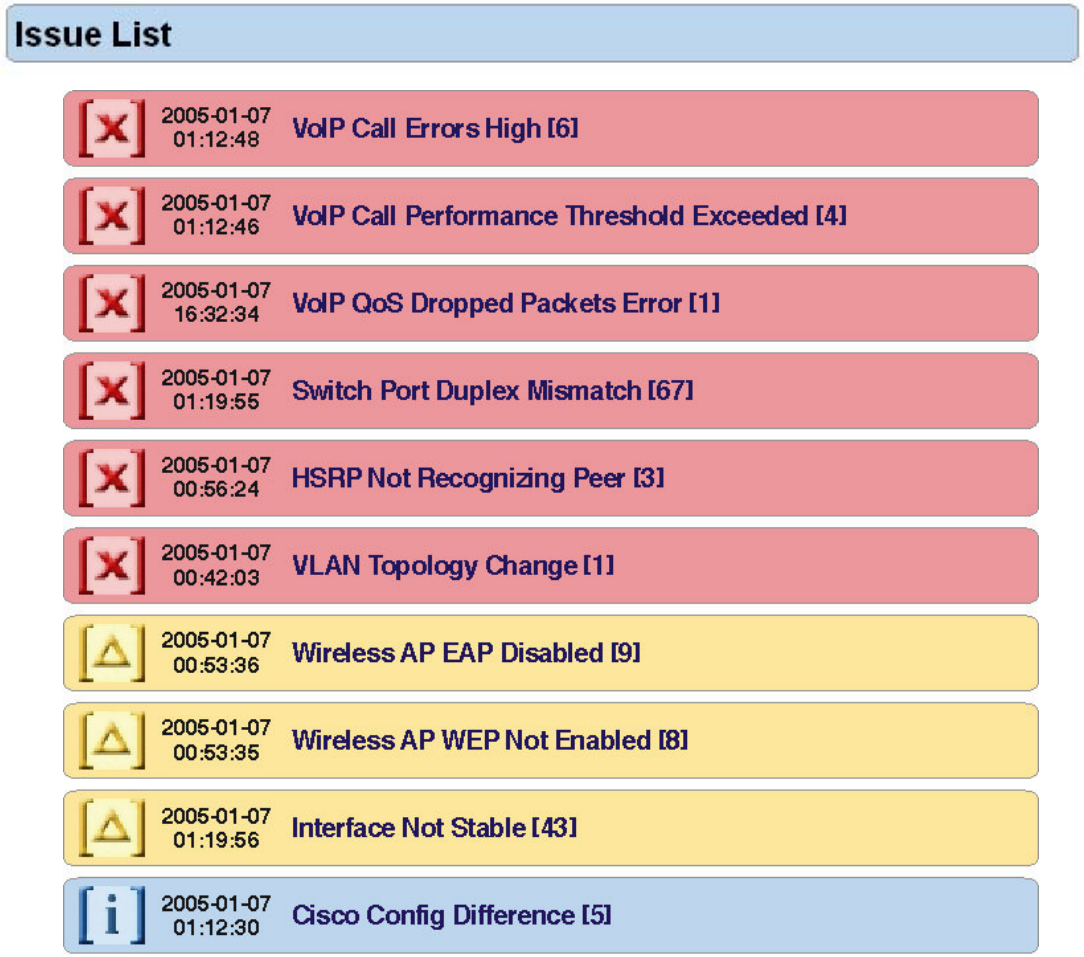

Premise - 2 Month Install is Good!

- Two months is fast for some systems
  - Customization (unique networks)
  - Training
- Bad for Return On Investment
- How long until results?
  - Days, weeks, or months?

Premise - My Network is Unique

"If your network is unique, it has unique problems."
- Mike O'Dell, former Chief Scientist at UUNET

- Good networks share common characteristics
  - Basic design, operational goals
  - Especially within an industries (e.g., financial, healthcare)
  - Uniqueness of addressing and applications
  - Uniqueness in level of redundancy

Premise - My Network is Unique

- Building block design
  - Avoid "Anything to Anything" connectivity
  - Known failure modes

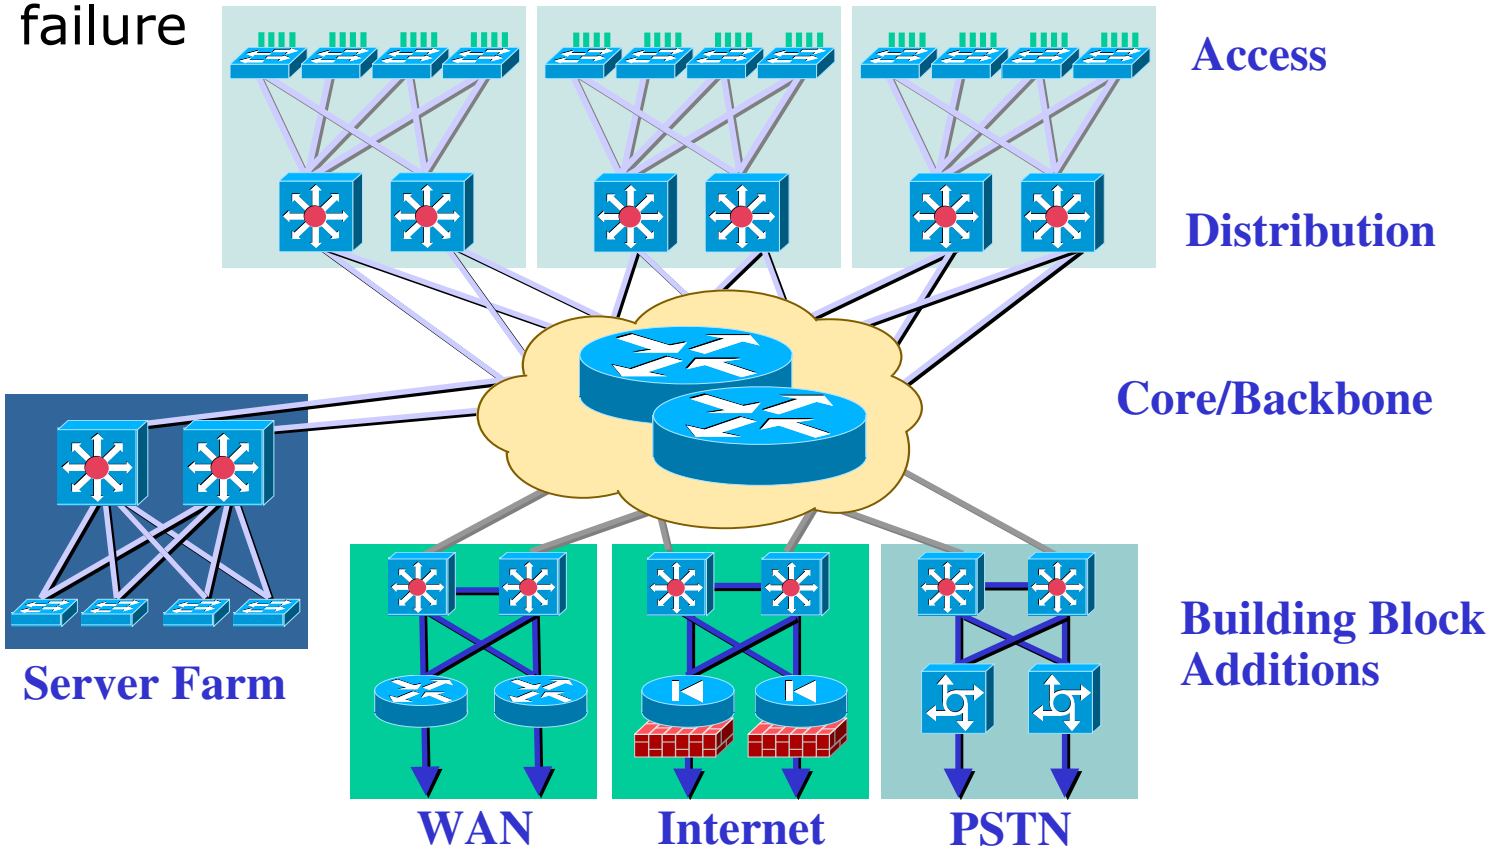

UI - Visualization

- Interesting display…
- Bowl of spaghetti
- May show high level structure: core, distribution, access layers
- Usefulness for troubleshooting?

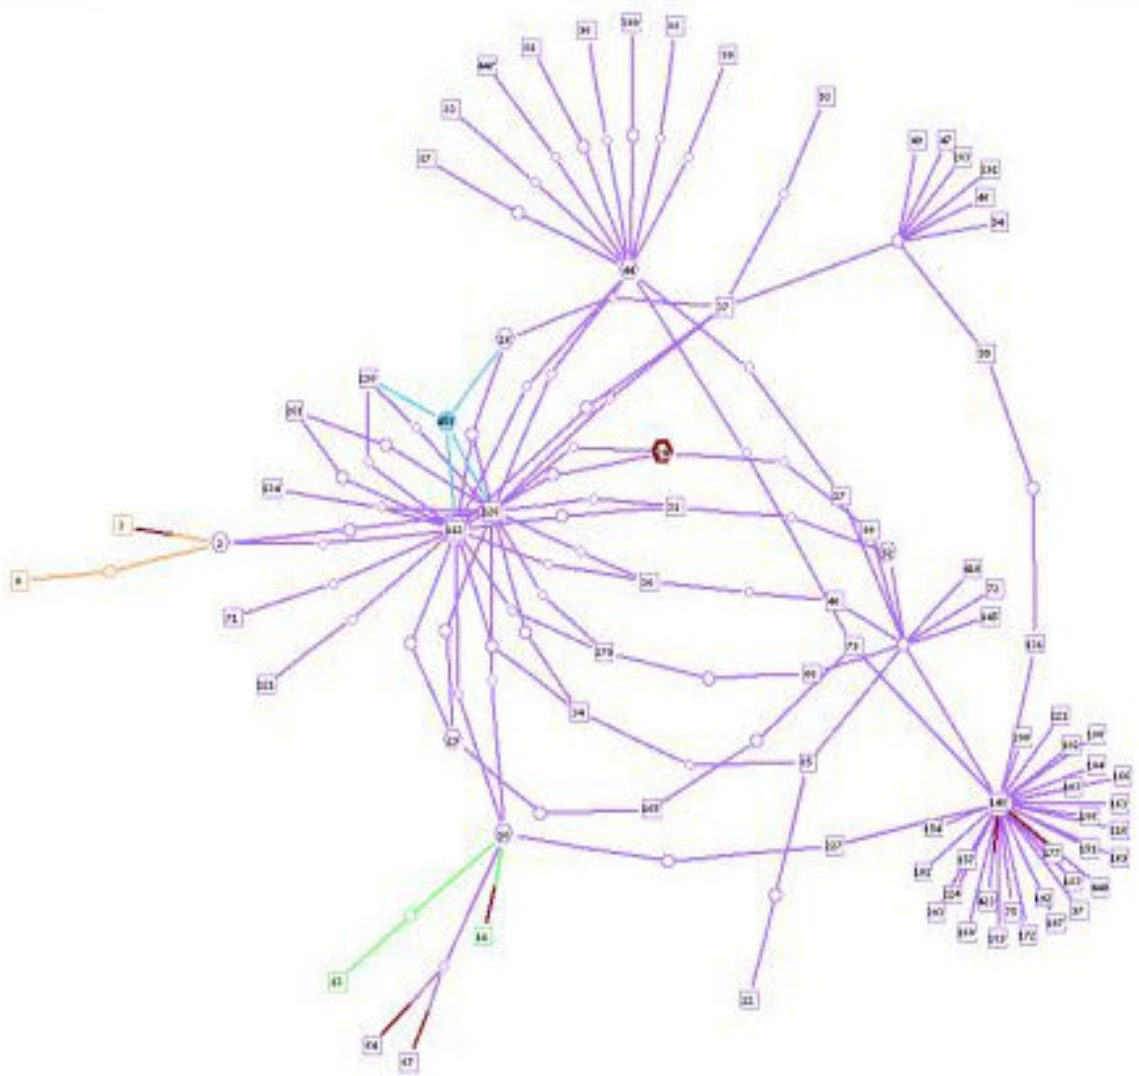

UI - Visualization

- Better example
- Making this useful...

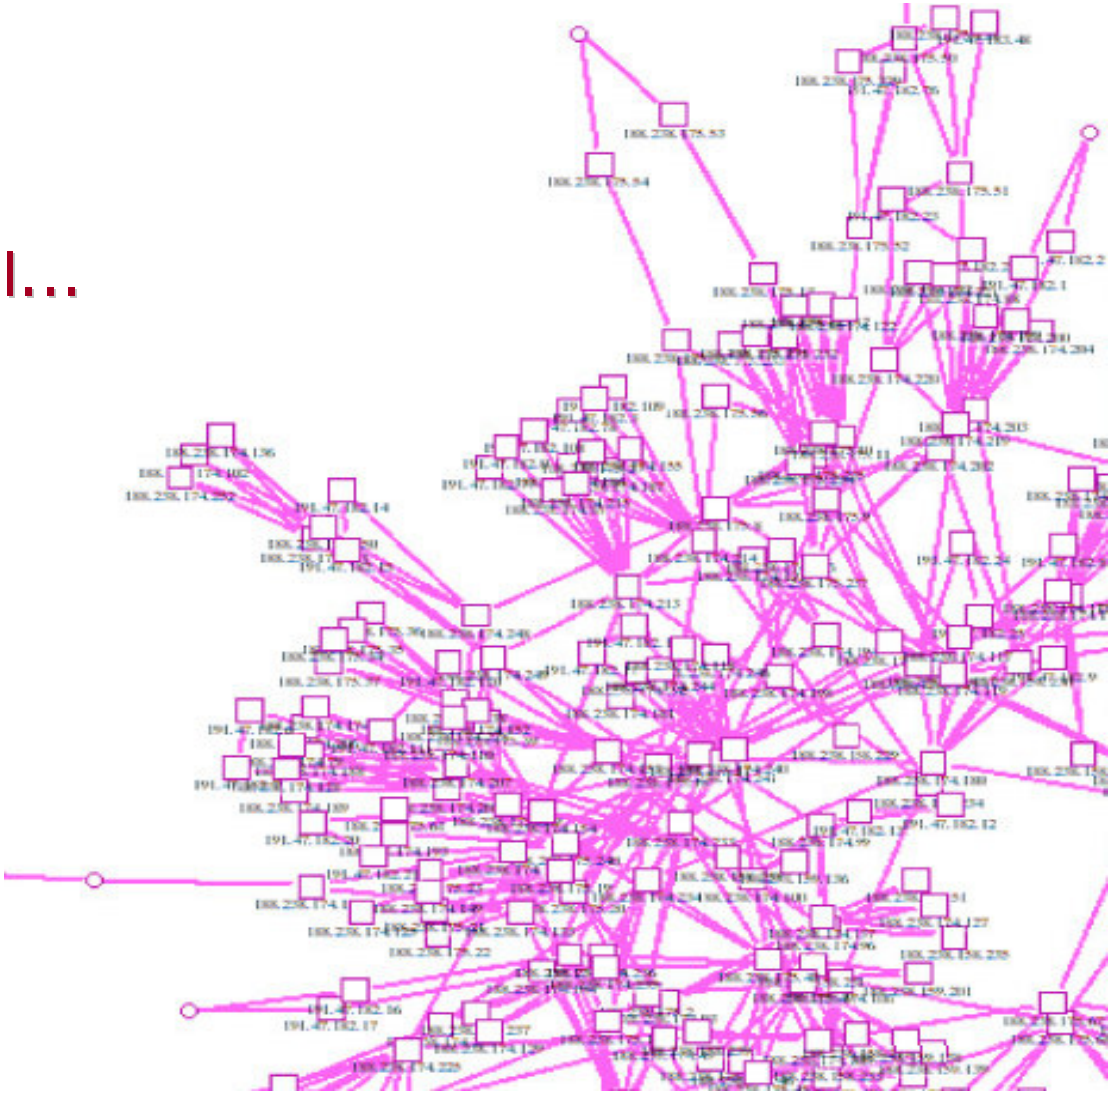

UI - Visualization

- Show the path of interest
- Troubleshooting aid
- Identify issues along the path

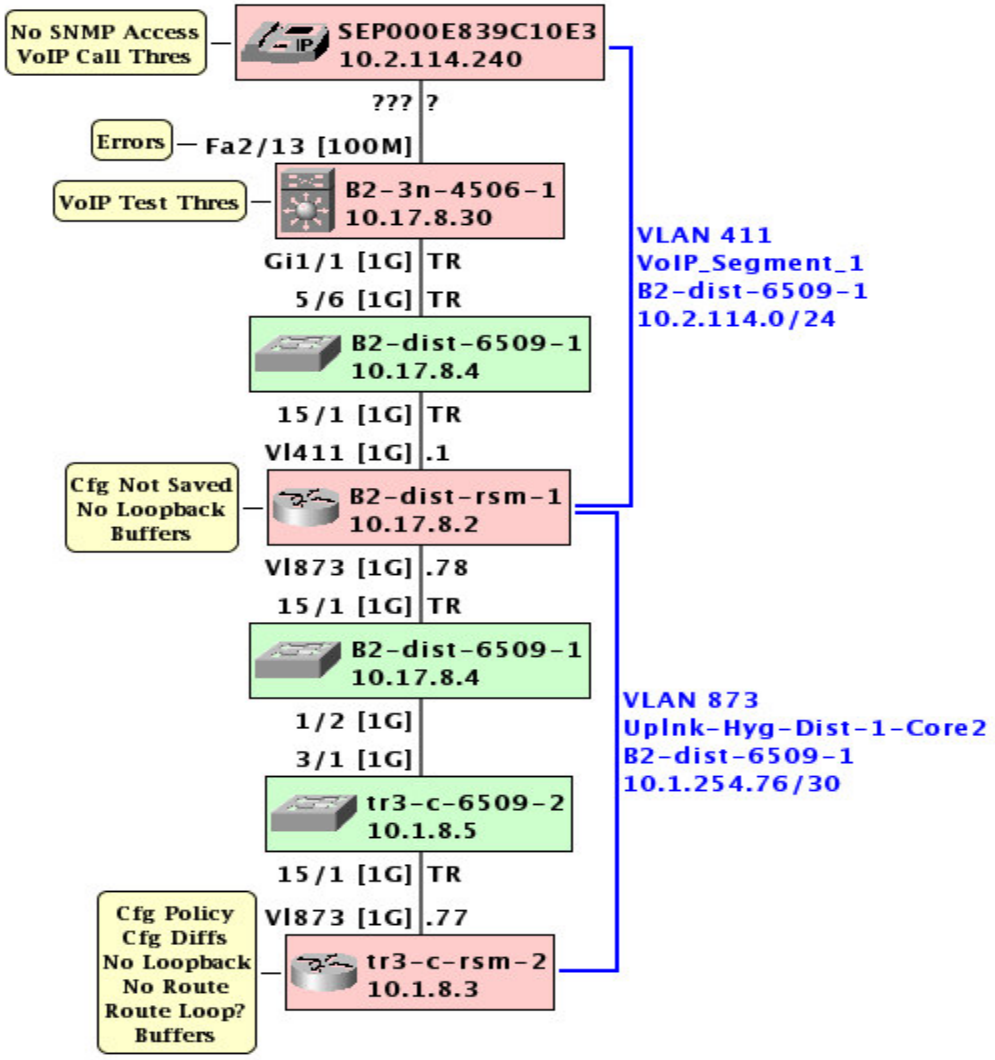

Premise - Using Managed Devices

- "Managed devices are too expensive" !?

- Reduce network staff load
  - Gain visibility into network operation
  - Remote administration
- Improve troubleshooting
  - Easier troubleshooting
  - Shorter time to repair
- ISPs learned this lesson
  - Modem pools
  - Automatic reset
  - Reduces staff requirements
  - Improved customer service

Data vs Information - Bridge Priority

- Data: bridge priority per switch

```
Switch-12> (enable) show spantree 1
   VLAN 1
   spanning-tree enabled
   spanning-tree type          IEEE
   Designated Root             00-10-0d-b1-78-00
   Designated Root Priority    32768
   Designated Root Cost        19
   Designated Root Port        2/3
   Root Max Age   20 sec    Hello Time 2  sec   Forward Delay 15 sec
   Bridge ID MAC ADDR          00-10-0d-b2-8c-00
   Bridge ID Priority          32768
   Bridge Max Age 20 sec    Hello Time 2  sec   Forward Delay 15 sec
```

Data vs Information - Bridge Priority

- Information: analysis of values WRT each other
- No events generated

VLANs

Rows 41-60 of 176

	VLAN ID	VLAN Name	Root Bridge	Count
41	110	1st_Floor	B2-dist-6509-1	7
42	120	2nd_Floor	b1-dist-6509-1	7
43	140	4rd_Floor	B2-dist-6509-1	7
44	411	VoIP_Segment_1	B2-dist-6509-1	6
45	900	Multimedia_SupportService	B2-dist-6509-1	6
46	901	Kiosk_IS	B2-dist-6509-1	6
47	902	FitnessCenter_vl902	t56-dist-6506-1	6
48	951	Daily_Grind	t56-dist-6506-2	6
49	1	default	tr3-c-6509-1	5
50	120	2nd_Floor	B2-dist-6509-1	5
51	130	3rd_Floor	b1-dist-6509-1	5
52	170	7th_Floor	t12-dist-6506-1	5
53	170	7th_Floor	b1-dist-6509-1	5

VLAN Root Details

VLAN ID :	1	Bridge Max Age :	2000
VLAN Name :	default	Bridge Hello Time :	200
Root Bridge :	tr3-c-6509-1	Bridge Fwd Delay :	1500
Root Priority :	8192	Top Changes :	0
Root Bridge ID :	0x20:00:00:0A:42:B0:B4:00		

VLAN Switches

Rows 1-5 of 5

	Device Name	VLAN Name	Priority	Bridge Address	Timers
1	rmt-sites-2950-1	default	49152	00:09:B7:F7:78:C1	OK
2	tr3-c-6509-2	default	16384	00:0A:42:B0:A4:00	OK
3	**tr3-c-6509-1**	**default**	**8192**	**00:0A:42:B0:B4:00**	**OK**
4	tr3-VoIP-2950-1	default	49153	00:0A:8A:89:E0:C1	OK
5	dist-ed-4503-1	default	32769	00:0C:CE:96:53:80	OK

CSV Data

UI - Strip Charts

- Strip-chart-itis
- How many interfaces can be monitored?
- What is the speed of these interfaces?
- Only useful for diagnostic purposes

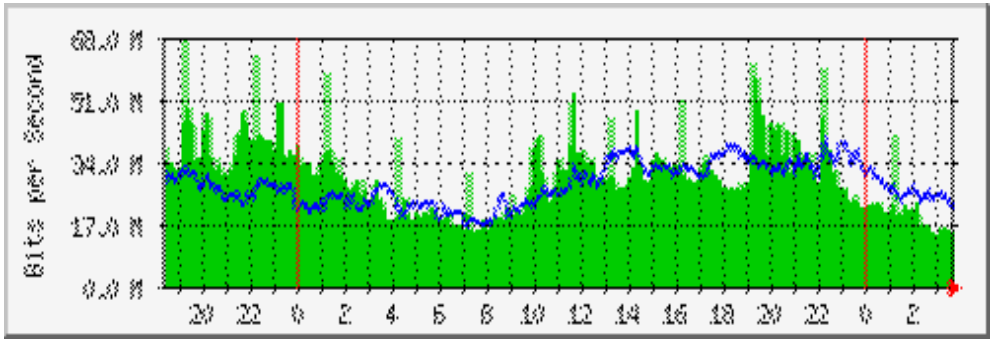

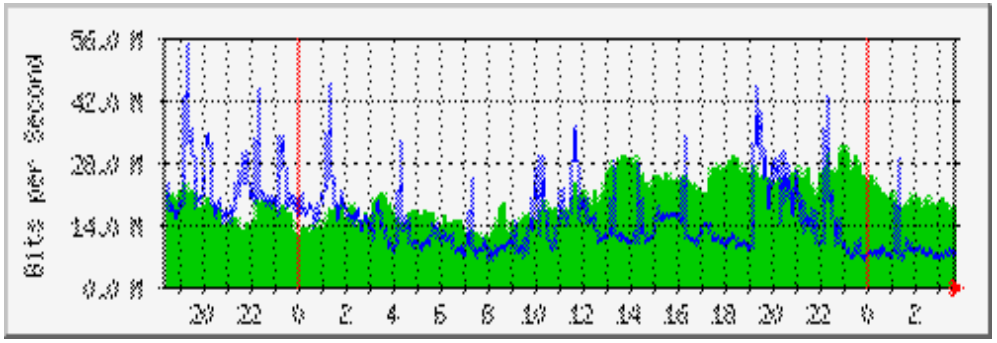

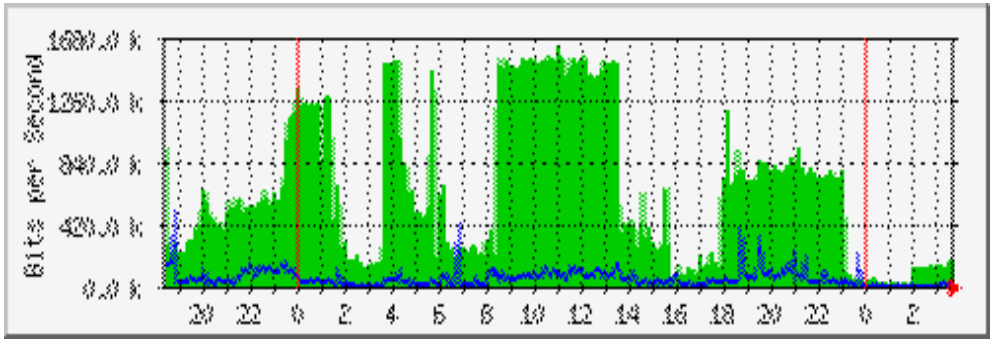

Polling - SNMP Packet Size

- Poorly implemented SNMP agents
- Maximum of 484 bytes on some devices
- Minimum packet size of 484 bytes!
- Question vendor's commitment to network management

Data vs Information - HSRP

- Data: HSRP configuration per router
- Information: active & standby router

HSRPs

Rows 1-20 of 107

	Virtual IP Address	Group Number	Count
1	10.1.8.1	30	2
2	10.1.114.1	41	2
3	10.1.161.1	161	2
4	10.1.217.1	217	1
5	10.1.254.90	80	2
6	10.1.254.98	81	2
7	10.1.254.241	72	2
8	10.4.8.1	30	2
9	10.4.9.1	9	2
10	10.4.27.1	27	2
11	10.4.90.1	90	2
12	10.4.99.1	99	2
13	10.4.207.1	207	2
14	10.4.213.1	213	2
15	10.4.215.1	207	2

HSRP Active Router Details

Virtual IP Address :	10.1.114.1	Auth :	cisco
Device Name :	tr3-c-rsm-1	Priority :	110
Interface :	VoIP Server Segment	Preempt Delay :	0
Group Number :	41	Configured Hello Time :	3000
State :	active	Configured Hold Time :	10000
Active Router :	10.1.114.2	Learned Hello Time :	3000
Standby Router :	10.1.114.3	Learned Hold Time :	10000
Virtual MAC Address :	00:00:0C:07:AC:2		

HSRP Members

Rows 1-2 of 2

	Device Name	Interface	State	Priority	Preempt Delay	Config Hello Time	Config Hold Time	Learn Hello Time	Learn Hold Time
1	tr3-c-rsm-1	Vl411 - VoIP Server Segment	active	110	0	3000	10000	3000	10000
2	tr3-c-rsm-2	Vl411 - VoIP Server Segment	standby	90	0	3000	10000	3000	10000

CSV Data

Data vs Information - HSRP

- HSRP groups containing only one router
- That's useful information!

HSRPs

Rows 1-20 of 107

	Virtual IP Address	Group Number	Count
1	10.1.8.1	30	2
2	10.1.114.1	41	2
3	10.1.161.1	161	2
4	10.1.217.1	217	1
5	10.1.254.90	80	2
6	10.1.254.98	81	2
7	10.1.254.241	72	2
8	10.4.8.1	30	2
9	10.4.9.1	9	2
10	10.4.27.1	27	2
11	10.4.90.1	90	2
12	10.4.99.1	99	2
13	10.4.207.1	207	2
14	10.4.213.1	213	2

HSRP Active Router Details

Virtual IP Address :	10.1.217.1	Auth :	cisco
Device Name :	tr3-c-rsm-1	Priority :	110
Interface :	B2 2ndFl Distant Ed Lab	Preempt Delay :	0
Group Number :	217	Configured Hello Time :	3000
State :	active	Configured Hold Time :	10000
Active Router :	10.1.217.2	Learned Hello Time :	0
Standby Router :	0.0.0.0	Learned Hold Time :	0
Virtual MAC Address :	00:00:0C:07:AC:D		

HSRP Members

Rows 1-1 of 1

	Device Name	Interface	State	Priority	Preempt Delay	Config Hello Time	Config Hold Time	Learn Hello Time	Learn Hold Time
1	tr3-c-rsm-1	Vl258 - B2 2ndFl Distant Ed Lab	active	110	0	3000	10000	0	0

CSV Data

Premise - Events are Sufficient

- Syslog & SNMP Traps
  - UDP based - packet may be lost
  - High volume - thousands per day
- Maintaining event filters
  - Constant maintenance
  - Like maintaining ACLs
- Many network problems don’t generate events
  - Duplex mismatch
  - Root bridge selection
- Correlating events
  - Ping failed: device down, or network down?
  - One fault may trigger many events

Network Configuration Access

- Unreliable, insecure access methods still used!
- TFTP
  - Transport is not secure
  - Unreliable in long delay paths (UDP based w/ timeouts)
- Telnet
  - Transport is not secure
  - TCP based, reliable transfer
- SSH/SCP
  - Transport is secure (SSHv2)
  - TCP based, reliable transfer
- Silliness:
  - "We're not allowed to use freeware." - Major financial firm, in reference to using SCP

Data vs Information - HSRP & VLAN

- Coordinating HSRP and VLAN root bridge

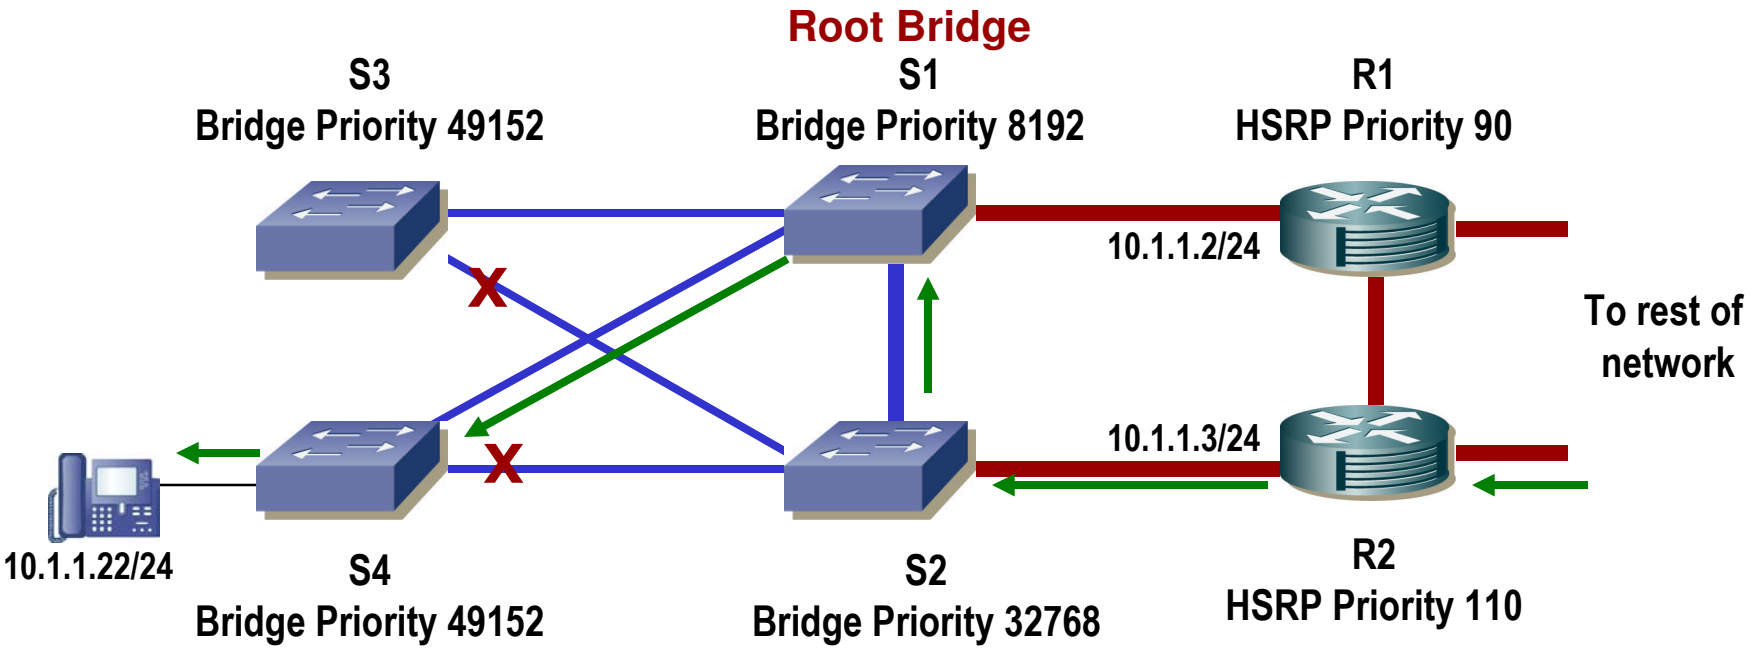

UI - Counts With No Basis

- What does this error graph tell you?

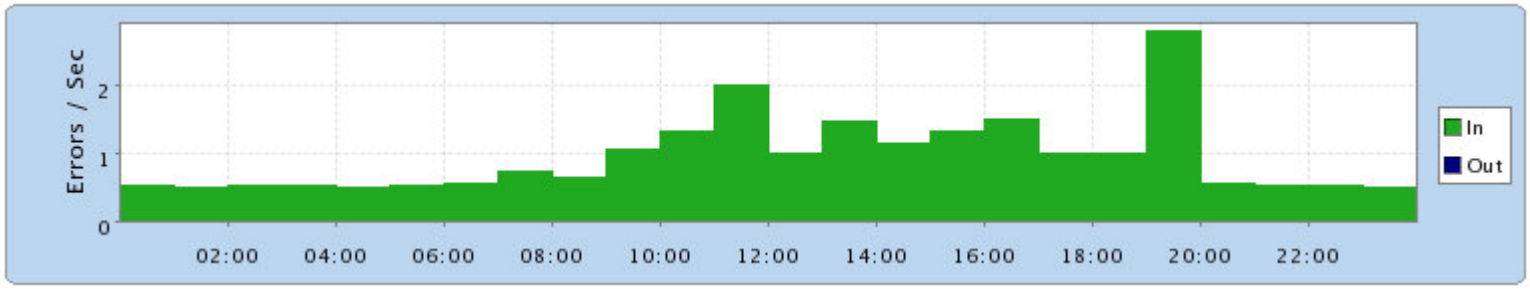

- How about this error graph?

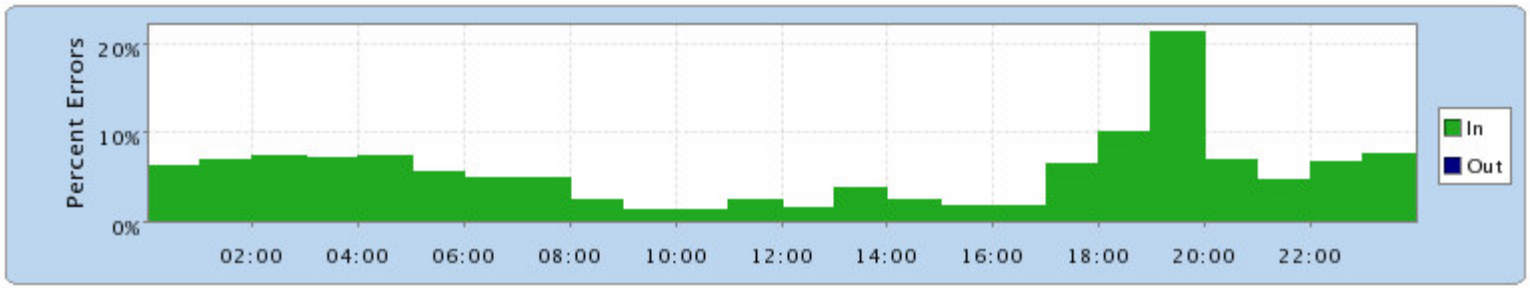

Polling - Lack of Standard Variables

- Every SNMP application has to deal with lack of standard variables

- CPU & memory
  - Support for multiple processors
  - Total memory, memory in use
  - Process table with CPU and memory stats
- Unique device ID
  - Recognizing each unique device when addresses change
  - Maintenance contract tracking (serial number)

    Cisco's 'show version':
    Processor board ID JAB04150AYP (161193301)

    SNMP says: 161193301

Data vs Information - ipOutNoRoutes

- Neat looking SNMP variable

ipOutNoRoutes OBJECT-TYPE
SYNTAX Counter
ACCESS read-only
STATUS mandatory
DESCRIPTION
"The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this `no-route' criterion. Note that this includes any datagrams which a host cannot route because all of its default routers are down."
::= { ip 12 }

- Possible routing failure indicator

Data vs Information - ipOutNoRoutes

- Key MIB phrase:

  "Note that this includes any datagrams which a host cannot route because all of its default routers are down."

  - ARP failures are included in Cisco implementation

- Determining problem origin
  - Need the source and destination addresses
  - Segregate ARP failures from routing failures

UI - Human or Machine?

- Manual scheduling of network discovery
  - Computers are good at scheduling; people aren't
  - Keeping the system in sync with the network
- Associating SNMP community string with devices
  - Unproductive busy-work for people
  - Computers can do the association (with some bad community string alarms)
- Only manage devices you know
  - Seed file of devices to manage
  - Finding rogue devices (Linksys wireless router; unmanaged hub)

Premise - I Need Realtime

- Realtime
  - Reactive; like firefighting
  - Tactical
  - DS3 to New York is down
- Proactive
  - Preventive; like fire alarms
  - Strategic & tactical
  - Duplex mismatch
  - HSRP with 1 router

Switch Port Duplex Mismatch [77]

Rows 1-20 of 77

	IP Address	Device Name	Interface		Total Packets	% Errors
1	10.19.8.7	b1-lab-3524-2	Fa0/19 - Tech Office	In Out	6,590,029 3,994,156	0.41 0
2	10.1.8.4	tr3-c-6509-1	2/1 - 10/100 utp ethernet (cat 3/5)	In Out	4,464,732 4,641,789	34.46 0

HSRP Not Recognizing Peer [32]

Rows 1-20 of 32

	Virtual IP	HSRP Group	Router IP	Router Name	Unknown Peer
1	10.1.217.1	217	172.27.22.5	tr3-c-rsm-1	Standby
2	10.17.2.1	91	10.17.8.2	b2-dist-rsm-1	Standby
3	10.17.8.1	30	10.17.8.2	b2-dist-rsm-1	Standby
4	10.17.16.1	90	10.17.8.2	b2-dist-rsm-1	Standby
5	10.17.32.1	100	10.17.8.2	b2-dist-rsm-1	Standby
6	10.17.48.1	110	10.17.8.2	b2-dist-rsm-1	Standby
7	10.17.64.1	120	10.17.8.2	b2-dist-rsm-1	Standby
8	10.17.80.1	130	10.17.8.2	b2-dist-rsm-1	Standby
9	10.17.96.1	140	10.17.8.2	b2-dist-rsm-1	Standby
10	10.17.112.1	150	10.17.8.2	b2-dist-rsm-1	Standby
11	10.17.128.1	160	10.17.8.2	b2-dist-rsm-1	Standby
12	10.17.144.1	170	10.17.8.2	b2-dist-rsm-1	Standby
13	10.17.160.1	180	10.17.8.2	b2-dist-rsm-1	Standby
14	10.17.176.1	190	10.17.8.2	b2-dist-rsm-1	Standby
15	10.17.224.1	66	10.17.8.2	b2-dist-rsm-1	Standby
16	10.17.250.1	250	10.17.8.2	b2-dist-rsm-1	Standby
17	10.19.2.1	91	10.19.8.2	b1-dist-rsm-1	Standby
18	10.19.8.1	30	10.19.8.2	b1-dist-rsm-1	Standby
19	10.19.10.1	40	10.19.8.2	b1-dist-rsm-1	Standby
20	10.19.16.1	90	10.19.8.2	b1-dist-rsm-1	Standby

Polling - Bandwidth Hog

- Implementation dependent
  - Potential to be very efficient
  - Example: 200 routers & switches w/ 5000 interfaces
    - 80Kbps input & 80Kbps output
    - Less than .1% of 100Mbps interface
- Build large packets to reduce round trips
- Collect only necessary data
- Be smart about collection intervals
  - Nyquist sampling theorem
  - Statistically significant number of samples
  - Tradeoff of frequency vs network load

Premise - We Have Plenty of Tools

- That's the problem! (too many tools)
  - Many point products
  - Seldom used well in production
- Keeping all the tools sharp
  - Separately staffed network management team
  - Trying to achieve integration between tools
- Multitude of user interfaces
  - Training!
  - Lack of common database

- Pick a small tool set and make them effective

Network Configuration Archiving

- Major vendor defaults to 10 copies in archive!
- Recommendations
  - Disk space is cheap - keep copies forever
  - Grab changes at least once a day (hourly is better)
  - tkdiff-style output for identifying changes

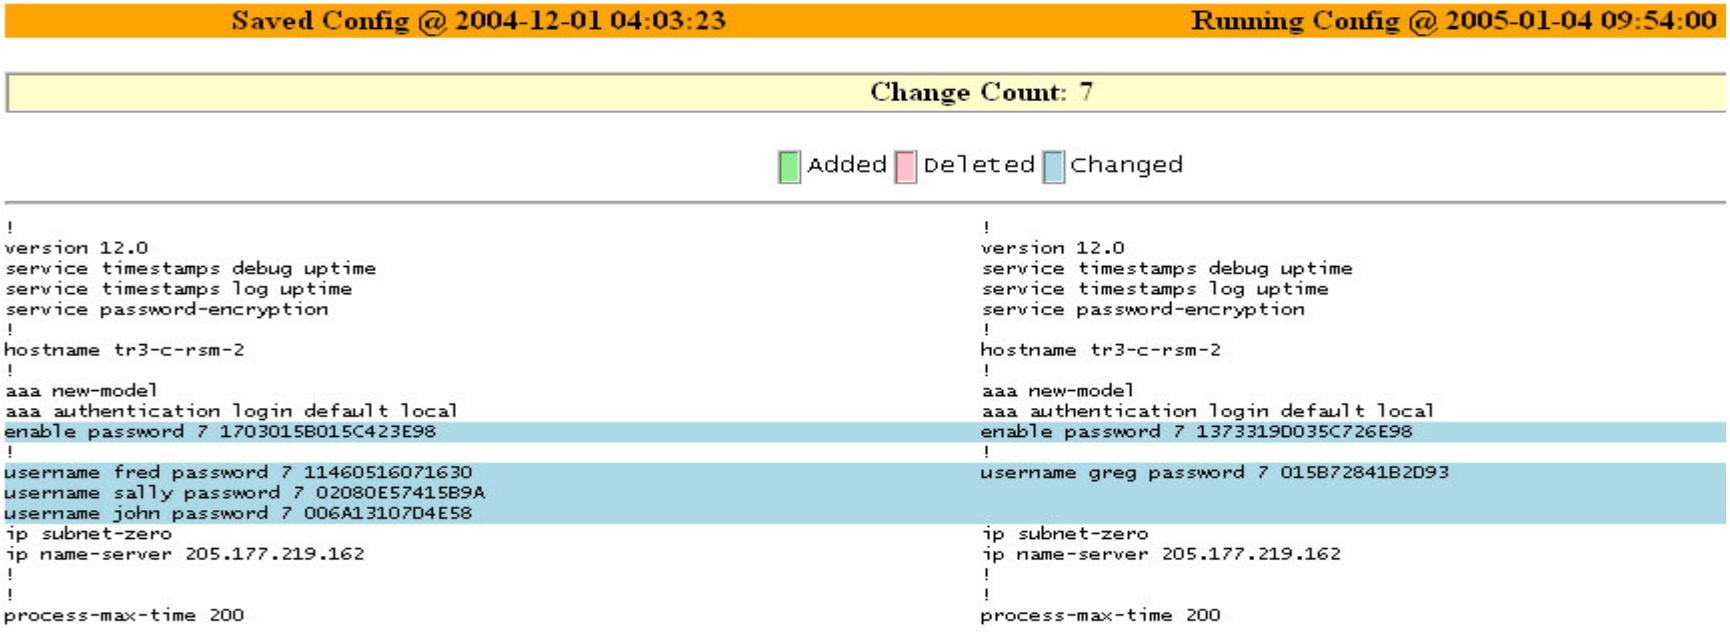

Saved Config @ 2004-12-01 04:03:23 | Running Config @ 2005-01-04 09:54:00

Change Count: 7

Added Deleted Changed

```
!                                                  !
version 12.0                                       version 12.0
service timestamps debug uptime                    service timestamps debug uptime
service timestamps log uptime                      service timestamps log uptime
service password-encryption                        service password-encryption
!                                                  !
hostname tr3-c-rsm-2                               hostname tr3-c-rsm-2
!                                                  !
aaa new-model                                      aaa new-model
aaa authentication login default local             aaa authentication login default local
enable password 7 1703015B015C423E98               enable password 7 1373319D035C726E98
!                                                  !
username fred password 7 11460516071630            username greg password 7 015B72841B2D93
username sally password 7 02080E57415B9A
username john password 7 006A13107D4E58
ip subnet-zero                                     ip subnet-zero
ip name-server 205.177.219.162                     ip name-server 205.177.219.162
!                                                  !
!                                                  !
process-max-time 200                               process-max-time 200
```

Polling - Large Tables

- Large vendor resorts tables on each GetNext!
  - Causes high CPU load
  - SNMP is low priority process
- GetBulk isn't significantly better

- Can't the table be cached until it is invalidated?

Data vs Information - icmpInTimeExcds

```
icmpInTimeExcds OBJECT-TYPE
                SYNTAX  Counter
                ACCESS  read-only
                STATUS  mandatory
                DESCRIPTION
                        "The number of ICMP Time Exceeded
                        messages received."
                ::= { icmp 4 }
```

- Possible routing loop
  - But who?
  - Need source & destination addresses

Network Configuration Policy

- Checking configurations against policy
  - Traditionally manual process
  - Seldom done
- Establish a configuration policy
  - Addressing (infrastructure & voice segregated)
  - Security
  - Routing design
  - Switching design
- Policy compliance
  - Was the design deployed as intended?
  - Tracking policy changes
- Tools now exist

Polling - Community String Indexing

- Ugly hack instead of fixing the MIB
- public@2 to access VLAN 2's bridge MIB data
  - MIBs that have only one instance
  - Vendor doesn't wait to fix MIB
- Pushes the problem into the network management app

Data vs Information - Reboots

- Switch rebooting
  - Cause?
  - 'show version' says "Corrupt PC"
- Probable Causes
  - Memory fault?
  - Heat?
  - IOS bug?

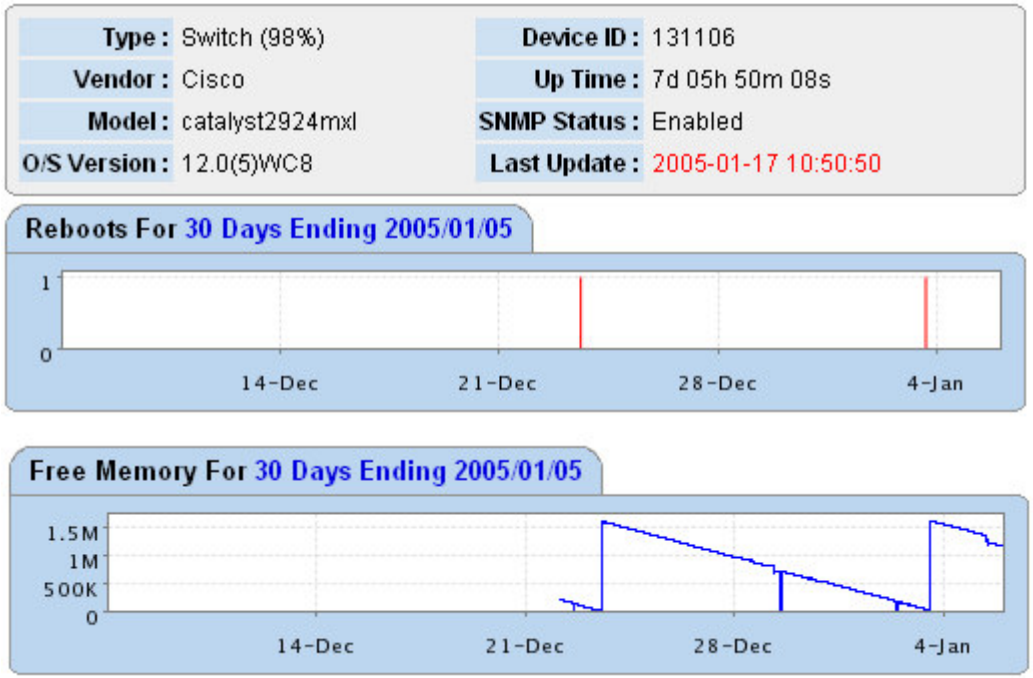

- Cause: IOS bug
  - Memory leak
  - 12 day reboot period
  - Need correlation of data to discern

Summary

- Many broken things
- Workarounds prevail at every step
- Combine data from multiple sources
  - Create information, not data
  - Useful displays are critical
- Breadth is necessary
  - Faults (events)
  - Performance
  - Accounting
  - Operational configuration
  - Configuration changes
  - Policy compliance

tcs@netcordia.com
www.netcordia.com/lisa2005