

Fun With Traces – Chapter 1

<http://www.skendric.com/seminar/>

Stuart Kendrick

Systems Engineer

Allen Institute for Brain Science

The Concept

Fun With Traces

I developed this seminar as a day-long Hands-On Lab, in which we use a puzzle-solving approach to analyzing real-world cases, learning protocols behave and practicing Wireshark techniques along the way. And generally, this is the format in which I deliver it, as a day-long class, during which we alternate between individual practice time and group discussion.

Alternatively, I can present this in more of a seminar-style, in which we review the cases together, brain-storm next steps, and proceed to the solutions, without actually practicing individually – less effective pedagogically, but can be delivered in a couple hours.

I try to slip a lot of lessons into this format, from refining the Problem Statement to diagramming IT environments to leveraging your Problem Management process for communicating risk. Some of the specific tips are Wireshark-specific, others apply to any protocol analyzer, and some aren't related to analyzers at all but are instead insights I've found useful for understanding how clients and servers communicate.

See <http://www.skendric.com/seminar> for related material, including *Homework* – takes you through basic Wireshark lessons.

Mechanics

Talk

- I encourage interactivity
- If you want to contribute, feel free to interrupt me
- Or raise your hand, and I'll call on you
- I'm good with either approach

Traces

- Grab a USB stick from the table up front

If you're reading this on-line, you may notice that some of the slides appear to be missing ... yes, I suppress the slides containing the 'answers' when I publish to my Web site.

This deck available at <http://www.skendric.com/seminar/>

Me

Multi-disciplinary IT trouble-shooter / Root Cause Analysis

<http://www.skendric.com>

sbk@cornella	<i>student</i>	1981
stuart@cpvax5 (Science Applications Int Corp)	<i>programmer</i>	1984
sbk@cornellc.cit.cornell.edu	<i>desktop / server</i>	1985
stuart.kendrick@med.cornell.edu	<i>server / network</i>	1991
skendric@fhcrc.org	<i>multidisciplinary</i>	1993
stuart.kendrick@isilon.com	<i>sustaining engineer</i>	2013
stuartk {at} alleninstitute dot org	<i>systems engineer</i>	2014

IT Architect | ITIL Problem Manager | Problem Analyst | Device Monitoring | Transport

Geeky Highlights

PL/1 on IBM mainframes	<i>Cornell University</i>	<i>Ithaca</i>	<i>1981</i>
FORTRAN on CRAY-1	<i>SAIC</i>	<i>San Diego</i>	<i>1984</i>
Terak, DisplayWriter, IBM PC, Macintosh	<i>Cornell University</i>	<i>Ithaca</i>	<i>1985</i>
Netware, Corvus Omninet, TCP-IP / IPX / AppleTalk	<i>Cornell University</i>	<i>Ithaca</i>	<i>1988</i>
AppleShare, QuickMail, Farallon, NRC, Cisco, Sniffers	<i>Cornell Medical College</i>	<i>Manhattan</i>	<i>1991</i>
Solaris, Windows, Linux, Perl, SNMP, Wireshark, Cisco, Fluke	<i>FHCRC</i>	<i>Seattle</i>	<i>1993</i>
OneFS: Authentication:Identity Mgmt:Authorization	<i>EMC Isilon</i>	<i>Seattle</i>	<i>2013</i>
Scientific application support	<i>Allen Institute for Brain Science</i>	<i>Seattle</i>	<i>2014</i>

Geek credentials: I missed punch-cards by one semester ... grew up on shared machines (IBM and Cray) ... my first network ran at 1Mb/s over Cat 2 (Corvus Omninet) carrying IPX + AppleTalk with IP encapsulated in both. I bored a vampire tap (once) ... my first analyzer was a Network General Toshiba 286 laptop ... and alpha versions of EtherPeek

What's a Protocol?

Ring, ring, ring

1. This is Stuart Kendrick of EMC Isilon, may I help you?
2. Hi Stuart, this is Alan, how are you?
3. I'm fine Alan, beautiful day here, how's the weather in Minneapolis?
4. Sunny and warm also. Hey Stuart, I have a question ...

That is a protocol. It contains standardized request/response pairs, e.g.

- Line 1 I identify myself
- Line 2 The caller identifies himself
- Line 3 I give the caller a hint of my status and capabilities
- Line 4 The caller also offers hints about his status ... then he gets down to business and issues a request to perform a specific task

A New York City implementation of a similar protocol:

Ring, ring, ring

1. Stuart Kendrick, EMC
2. Hi Stuart, this is Alan, how are you?
3. Whaddaya want?
4. What the !@#\$ are you doing for me on the Cerner Escalation ...

ProTip

Protocol designers employ a finite set of design patterns.

After you've seen a bunch of protocols, you'll start to recognize those patterns.

Then, you'll sit down to analyze a new protocol ... and be able to pick it up rapidly, because you recognize the request/response pairs as merely variations on some other protocol which you already know.

Common elements of protocols

- Request/Response pairs

- Status codes (succeed / fail / more-processing-needed / etc.)

- Sequence numbers (permits ordering)

- Identifiers (uniquely identify the conversation / request / etc.)

- ...

Case Studies

Case 1	Many Applications Crash	BlueHeat
Case 2	Loading a Home Page	What Takes So Long?
Case 3	VMWare Cannot Mount SAN	The Router is Broken
Case 4	HL7 Transfers Interrupted	eGate Eccentricity
Case 5	Account Lockouts	Why OS X?

Case 1

Many Applications Crash

BlueHeat

Case 1: Background

This is the last week in November 2005. Earlier this year, we bought a mass storage device – a BlueArc Titan NAS head named *Indigo* sitting in front of 14 TB of Fibre Channel, SATA, and ATA attached disk trays. We have been migrating home + shared directories for two divisions (~1200 staff) from a flock of aging DAS-equipped file servers onto *Indigo*, along with scratch space for the MIS group.

The experience has been rocky. Starting in June, an OS memory leak caused key processes to hang and sometimes even head freezes, both requiring reboots to fix. A controller fried, requiring emergency downtime for replacement. A controller firmware bug mangled a volume, leading to data loss. We have been applying hot fixes, firmware upgrades, and OS upgrades every few weeks. Starting in August, users began reporting crashing applications – notably Outlook, although Word and Excel and other applications hang as well, intermittently – some days are fine, some days are bad. The MIS group’s Tidal jobs fail regularly.

Backups are slow and sometimes don’t complete – we aren’t meeting our 24 hour Recovery Point Objective, and we have no confidence that we can meet our 48 hour Recovery Time Objective. *Sometimes even simple file copies are slow!*

Case 1: More Background

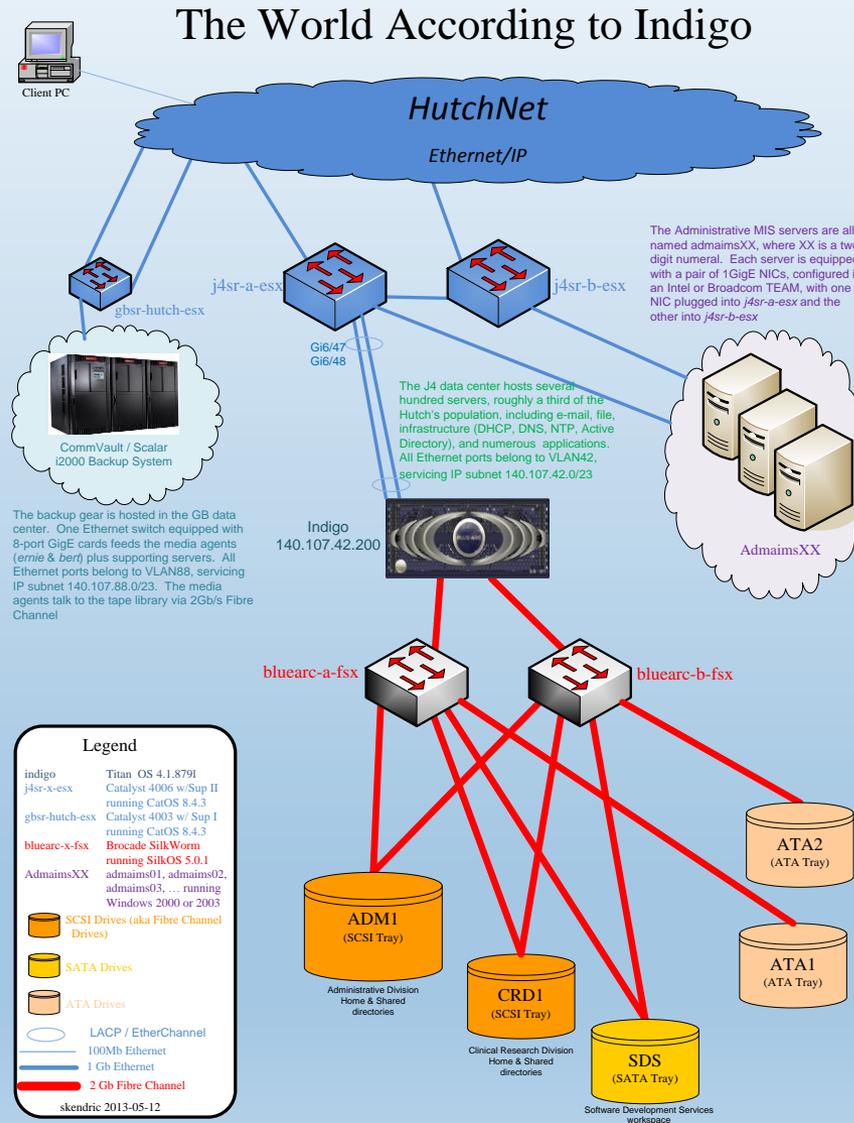
The storage team was convinced that antivirus scanning was causing the application crashes and has worked with BlueArc for months to resolve this, finally disabling AV over Thanksgiving. However, the intermittent application crashes continued this week.

The local BlueArc team visited a few days ago and identified the Catalyst 4000 Ethernet switches as the likely culprits: *“The Catalyst 4003 servicing the backup systems dates to 1998; the Catalyst 4006 servicing the Titan itself dates to 2000 – they are getting overwhelmed by traffic.”*

The remaining ~1500 users who have not migrated to *Indigo* are watching with dismay – currently, they are unaffected, scattered as they are between small NetApp NAS heads and a flock of aging file servers.

Management has made every Sunday night in December available to you for *Indigo* downtime – just ask.

Case 1: Many Applications Crash



Case 1: Problem Statement

Initial Problem Statement

The switch is overloaded and is dropping packets.

Improved Problem Statement

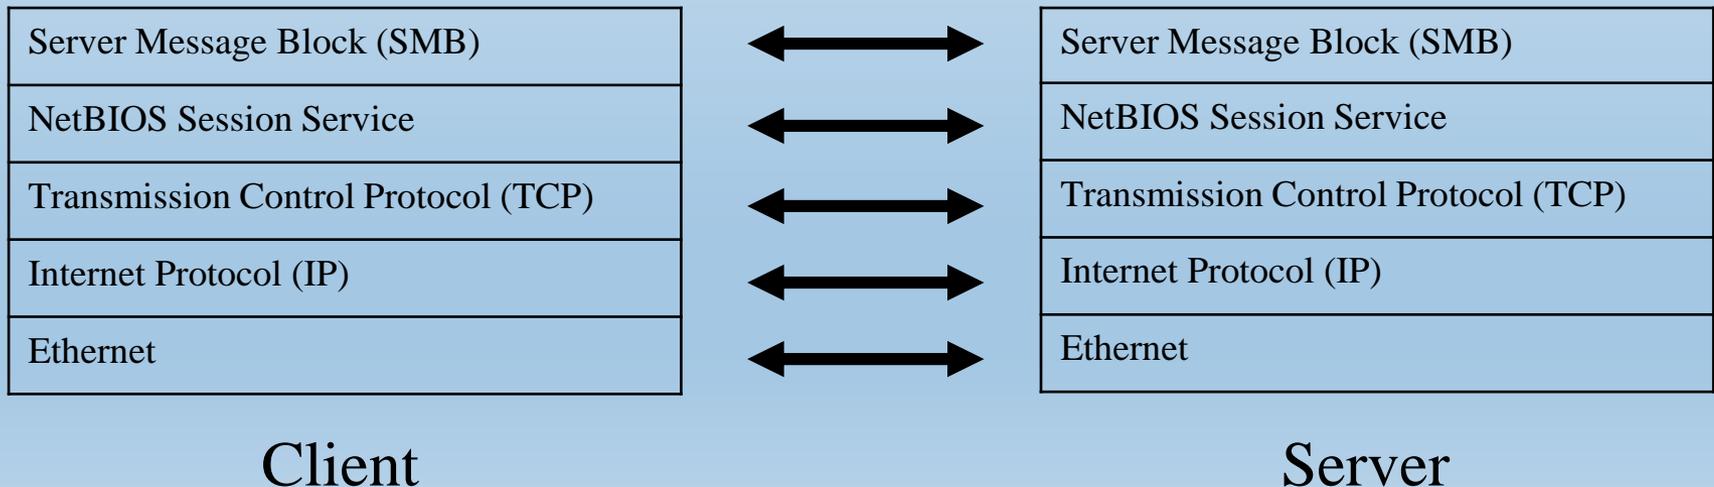
File copies are intermittently slow.

Buff it up a little more

Normally, a 10MB file copies in less than two seconds; but sometimes, the same file requires a minute or more to copy.

Case 1: Traces

- Glance at Baseline
 - Identify Request / Response pairs
 - View, Time Display Format, DeltaT
 - Sort by DeltaT
 - Sort by No.
- Glance at Slow trace
 - Sort by DeltaT
 - Who is introducing the delay?
- Search SMB documentation for SMB Echo



Case 1: Tell the Story

Who will recap for us?

Case 1: Potholes

- Under-powered gear: Indigo-side sniffer overwhelmed
- No Fibre Channel packet capture gear
- Reduced to a complex, tedious fishing expedition, looking for correlation amongst ~a thousand parameters ... took one of our lead analysts several weeks to find the key item
- And during that search, he found plenty of false correlation ... fortunately, he was skilled enough to resist the effect

Correlation provides clues, often misleading ones

In 1978, sports reporter and columnist Leonard Koppett mocked the causation-correlation confusion by wryly suggesting that Super Bowl outcomes could predict the stock market. It backfired: Not only did people believe him, but it worked -- with frightful frequency. If one of the 16 original National Football League teams -- those in existence before the NFL's 1966 merger with the American Football League -- won the Super Bowl, the stock market would close higher that following year than it did on the preceding Dec. 31. If a former AFL team won, it would go down. From 1967 to 1978, Koppett's system went 12 for 12; up through 1997, it boasted a 95 percent success rate. It stumbled in 1998 and 1999, when AFL alums the Denver Broncos won and the market went up.

Case 1: Tips

Packet trace quickly narrowed the fault domain

We gathered multiple traces, from multiple clients, on multiple days, increasing our confidence

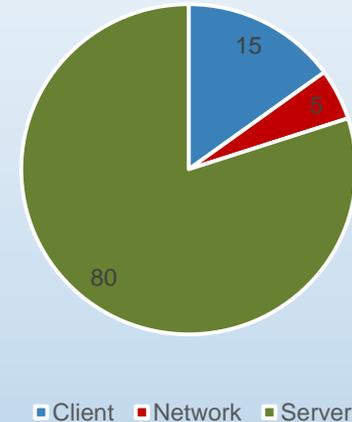
Sort by DeltaT: are there any big stalls?

Of course, this requires a focused, well-filtered trace, and if you have that, you're covered a lot of ground already, so this is a bit of an ingenuous tip

Are you still there?

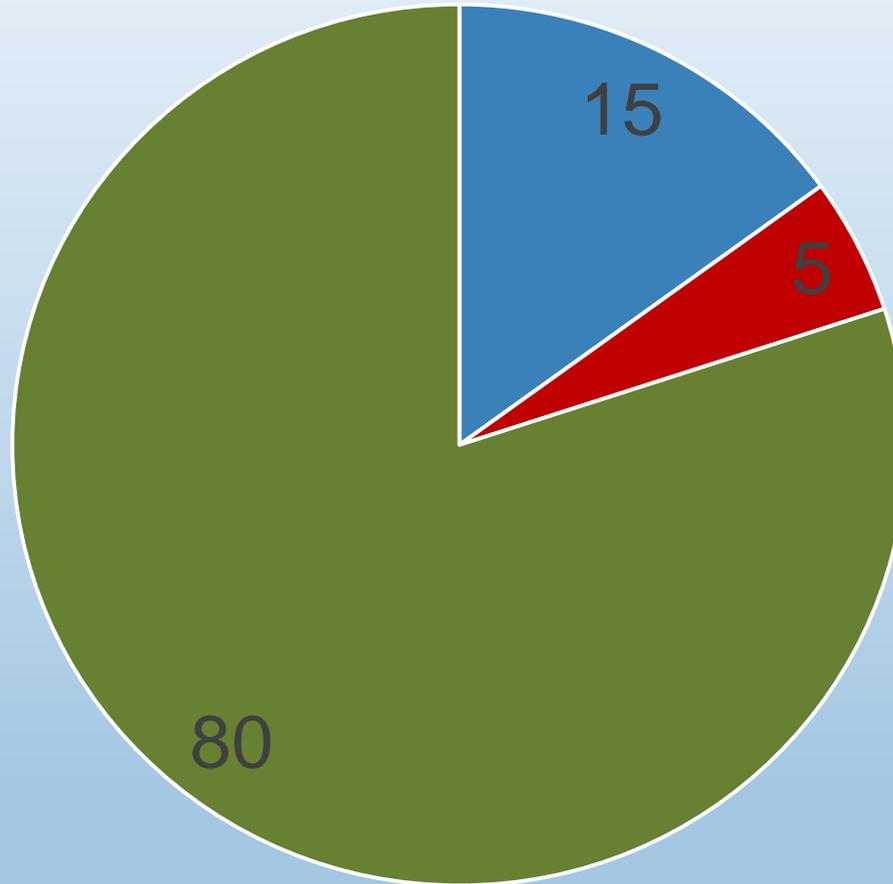
Protocols, and applications, tend to incorporate ways to 'ping' the other side to verify aliveness and even functionality: *SMB Echo*

Client - Network - Server Pie



Credible explanations grow from the combined testimony of three more or less independent, mutually reinforcing sources -- explanatory theory, empirical evidence, and rejection of competing alternative explanations. --Edward Tufte

Case 1: Client – Network – Server Pie



■ Client ■ Network ■ Server

Case 1: ProTips

- Compare a ‘good’ trace with a ‘bad’ trace
- Sort by DeltaT to find obscenely large pauses
- Watch for the Request/Response design pattern
- Look for stalls between the Request and the Response
- Notice that protocols implement ‘Are you there?’ mechanisms, e.g. SMB Echo
 - What might we have concluded if *Indigo* had **not** responded to the SMB Echo?
- Notice Wireshark Layout: Edit..., Preferences..., User Interface..., Layout
- Notice the *Resize All Columns* button
- Deep Lesson: Context is essential to pcap analysis: *Draw the Diagram, Tell the Story*

Case 1: Problems

Unable to capture Definitive Diagnostic Data

This is a Rapid Problem Resolution term: we did not have the tools to instrument the path from client to server: in particular, we did not have the gear to capture at *Indigo* nor inside the Fibre Channel network, and we had limited visibility into the Fibre Channel environment in general.

Propose

Add this to the list of risks tracked by Problem Management

Insufficient tools and access to diagnose client/server problems in the J4 data center

Case 2

Loading a Home Page

What Takes So Long?

Case 2: Tell the Story

My Client loads <http://www.emc.com> It takes 30 seconds. Why?

There isn't much pathology here; I'm just using this simple case to illustrate more Wireshark techniques.

Still, can we answer the question: why does my Client take 30 seconds to load this page?

Case 2: Filtering

I only need to look at a handful of packets to acquire a profound insight ...
But of course, finding that **particular** handful is hard...

Open EMC-Home-Page.pcapng

`ip.addr==192.168.75.20 and tcp.port==80`

Look for Request/Response pairs

Do I care about the TCP whinings?

Statistics..., Conversations ... How many TCP conversations does this trace contain?

Profile: Simple + TCP Stream

Open Frame 24

How does Wireshark know this frame is destined to TCP Port 80?

What does [Stream index: 1] imply?

Streams, aka Conversations, are defined by TCP Port pairs

See Frames 56-60 for more Conversations

Filter on Stream 1

IMHO: Filtering is the single most important concept in protocol analysis

Case 2: Columns, Profiles, Filters

OK, lots of Wireshark tips in the last few minutes; let's review a few

Change your Columns: Edit..., Preferences..., User Interface..., Columns

Add a new column ... Check-out Custom

Create Profiles

Create Filters

Review Filter Expressions section of

C:\Users\Stuart\AppData\Roaming\Wireshark\profiles\Simple + TCP Stream\preferences

Case 3

VMWare Cannot Reach

Storage

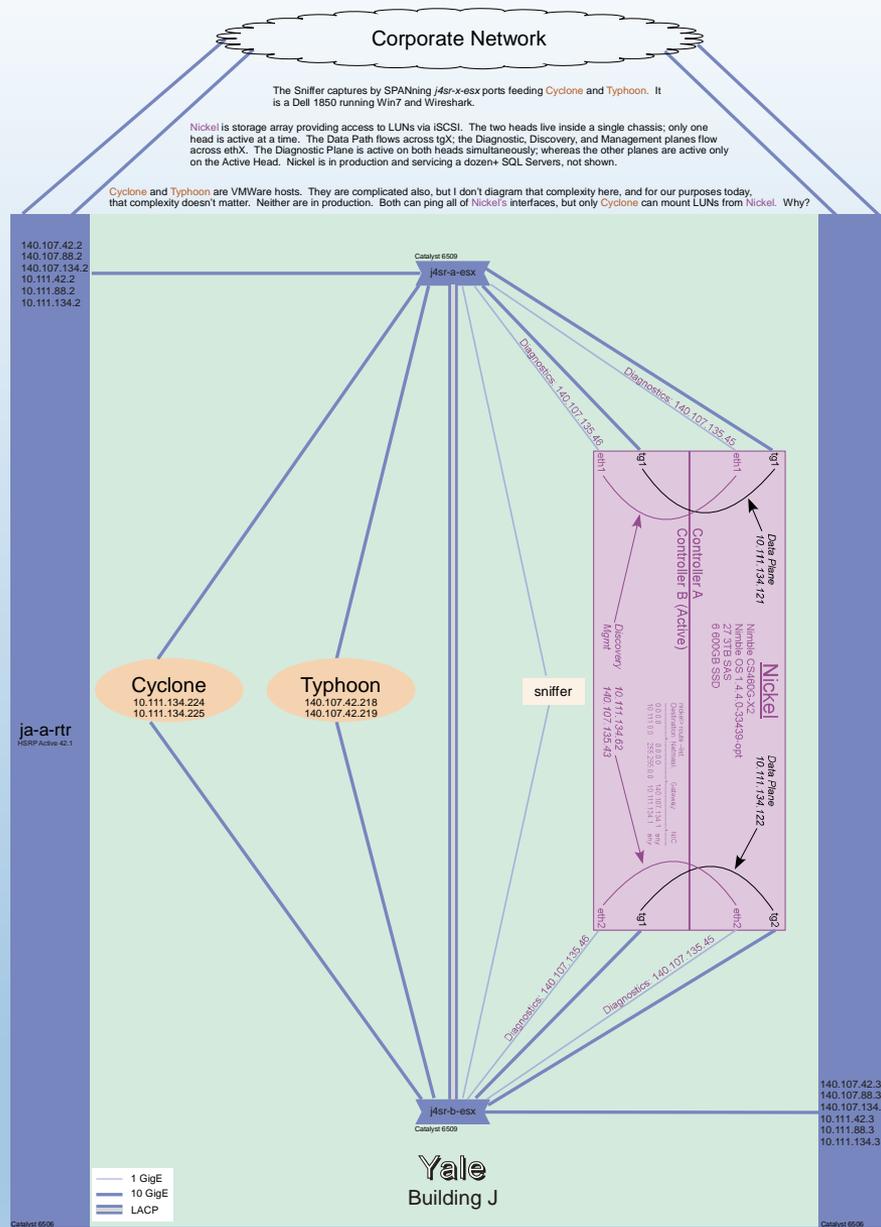
The Router is Broken

Case 3: Background

The storage folks have brought up their spiffy new SAN, a Nimble iSCSI box called *Nickel*. They are taking a gamble with it – Nimble was a new manufacturer at the time, just a few years old. During the eval and Proof of Concept phases, they found all sorts of immaturity. Still, the price was right, so they went with it. The box was intended to service both our VMWare farm (~1000 VMs distributed across 7 hosts) and a dozen SQL Servers (running PeopleSoft).

The VMWare farm moved first, and the results were spectacularly good (well, the previous SAN was overloaded, so just about anything would look good). But the SQL Servers had a different experience – the dev boxes did fine, but when production arrived, problems soared (performance and SQL crashes). After months of wide-spread disruption and slowness, the storage admins fixed a misconfiguration on *Nickel*, loaded a patch for a particular bug from Nimble, and life improved substantially.

The VMWare folks want to upgrade from VMWare 4 to VMWare 5 and refresh hardware as well. They have heated up two new boxes, *Cyclone* and *Typhoon*. *Cyclone* can mount LUNs off *Nickel* just fine, but *Typhoon* cannot. VMWare offers no error message – but the configured LUNs remain greyed out in the GUI. VMWare and Nimble tech support say everything is fine from their point of view and suggest that the router is broken, because *Cyclone* shares the same subnet as *Nickel*, but *Typhoon* is on a different subnet. Pings get through just fine and other hosts located in this data center and in other data centers mount iSCSI LUNs off *Nickel* (and other SANs) through this router just fine.



Case 3: Problem Statement

Initial Problem Statement

The router is blocking iSCSI

Improved Problem Statement

Typhoon cannot mount iSCSI LUNs from *Nickel*

Can you think of a better Problem Statement?

Case 3: Traces

A network admin and a VMWare admin collaborate

They first SPAN the *j4sr-x-esx* ports feeding *Cyclone* and capture a successful LUN mount into two trace files (they configure Wireshark to write two traces files – they don't know about Wireshark's ability to capture on multiple interfaces and write the results to a single trace file.)

And they repeat the process with *Typhoon*, capturing a failed LUN mount

Cyclone-Capture-on-j4sr-a-esx-Port merged with Cyclone-Capture-on-j4sr-b-esx-Port =
Cyclone-Both-Ports-Merged

Typhoon-Capture-on-j4sr-a-esx-Port merged with Typhoon-Capture-on-j4sr-b-esx-Port =
Typhoon-Both-Ports-Merged

[BTW: *Cyclone* and *Typhoon* are actually the same box, with IP addresses reconfigured. For our purposes today, this detail does not matter.]

Case 3: Your Story

How would you tell this story?

Case 4

HL7 Transfers Interrupted

eGate Eccentricity

Case 4: Background

Two servers, *Minnie* and *Simba*, exchange updates every hour, keeping each other apprised of changes made by their respective user bases. They each run *eGate*, an application gateway which specializes in exchanging Admit / Discharge / Transfer (ADT) information between medical IT systems, often between hospitals, allowing each side to transform the incoming data in ways which suit its internal systems. It speaks the HL7 protocol, a popular protocol for over-the-network data exchanges amongst healthcare systems.

These data exchanges sometimes fail. When that happens, a monitoring process sends e-mail to the on-call DataBase Admin, who can intervene to retry the transfer.

Historically, the DBA has ignored the message, relying on the next hourly data transfer event to succeed (this reliably executes both current updates and previously failed updates, i.e. no data is lost). However, the team is becoming concerned, as the frequency has been gradually increasing – last year, the frequency was once every few weeks; more recently, several times a day.

There are times when the user base would rather not wait two or more hours to see the latest data. Management has heard their concern, and now the on-call DBA must respond to the alarms, 7x24x365. Sometimes, this is happening several times a night.

Case 4: More Background

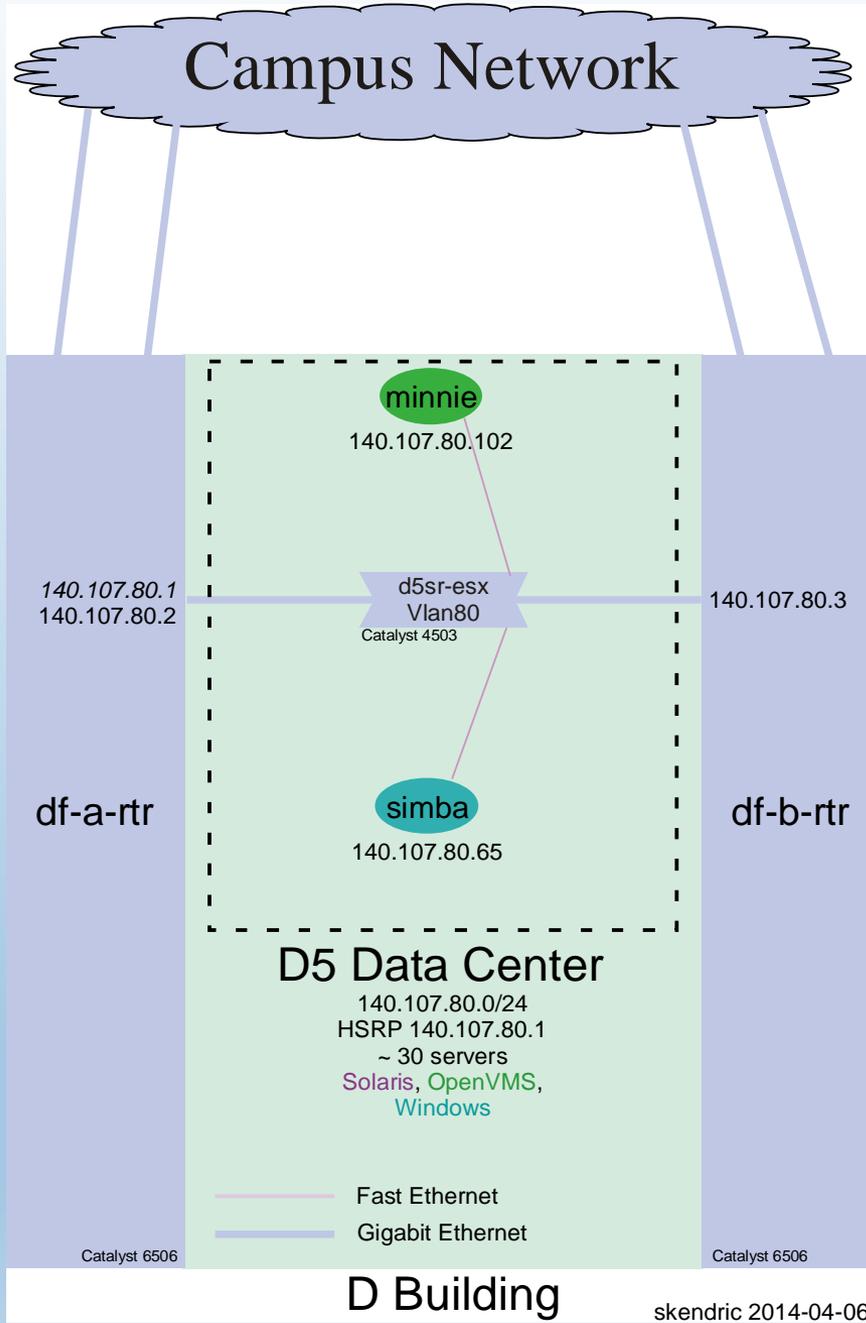
Management is concerned about the sleepless DBA issue; but they want you working on bigger problems and so ask you to ignore this.

You're sympathetic to your DBA buds – getting woken up each night is no fun.

One of the DBAs installed Ethereal on [Simba](#) years ago.

So, you show Mike how to set up a rolling capture; he understands how important it is to record event times precisely; and you tell that you'll spend an hour max analyzing whatever he finds.

A few days later, you get a handful of traces plus time estimates on when the transfer failed.



Case 4: Problem Statement

Initial Problem Statement

eGate transfers intermittently fail with ‘A Network Error has occurred’

Let’s look at one of them: *eGate-Fail-2011-04-17-0607.pcapng*

Xfer failed somewhere around 20:56

Filter on `tcp.stream==0`

Case 4: Tell the Story

How would you tell the story?

Case 5

Account Lockouts

Why OS X?

Case 5: Background

Historically, the password lockout threshold on the Active Directory controllers has been set to twenty (20). Yup, that's right, you get to type your password wrong 20 times in a row before the DCs lock your account.

The Security Office decides that this is too generous and draws the line in the sand at five (5).

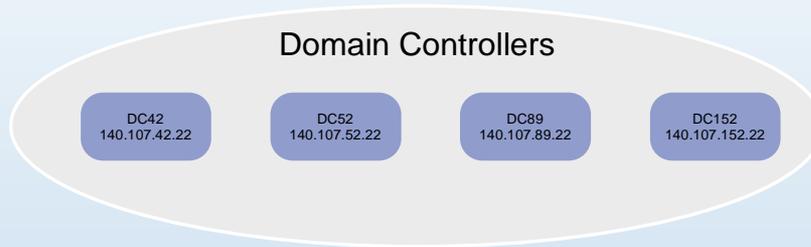
During Sunday night's change window, the AD administrator drops the threshold to five.

Monday morning, the HelpDesk reports an increase in Account Lockout calls. This continues all week and over the weekend. Next Monday morning, the Incident Commander, tired of being pestered, declares an Incident. You have been assigned to the Incident Response Team.

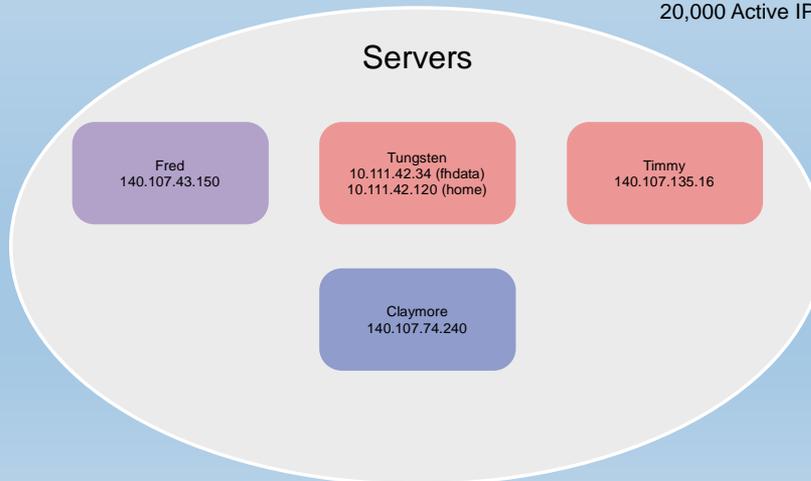
During the briefing, HelpDesk staff say that the problem seems to be isolated to OS X clients. On questioning, staff agree that they don't always ask what operating system the user is running, and that some callers have definitely been running Windows. One staff member notes that she has received calls from Linux users. However, the staff insist that they 'feel that' OS X is particularly affected.

End-User Community: 80% Windows, 15% OS X, 5% Linux

Account Lockouts



Not shown
Campus Network
8000 Active Ethernet ports
20,000 Active IP addresses



Case 5: Problem Statement

Initial Problem Statement

Accounts Lockouts have increased in frequency from a handful a day to several dozen a day (numbers not precise)

Case 5: Tell the Story

How would you tell the story?

SHOULDN'T BE HARD

|<

< PREV

RANDOM

NEXT >

>|



Wrap-Up

Questions, Comments, Complaints?

Thank you!

On-Line Resources

[Rapid Problem Resolution](#) by Paul Offord

LinkedIn [Protocol Analysis & Troubleshooting Group](#)

Old Comm Guy <http://www.loveytool.com>

Trouble-shooting & Training Outfits

James Baxter

<http://www.packetiq.com>

Tony Fortunato

<http://www.thetechfirm.com>

Chris Greer

<http://www.packetpioneer.com>

Paul Offord

<http://www.advance7.com>

Mike Pennacchi

<http://www.nps-llc.com>

Ray Tompkins

<http://www.gearbit.com>

...

Based Here (will travel for \$\$)

Daytona Beach, FL

Toronto, Canada

Central/South America

London (international)

Seattle, WA

Austin, TX

Conferences

Sharkfest

<http://sharkfest.wireshark.org>

San Francisco, CA

Follow-up

stuart.kendrick.sea {at} gee mail dot com

This deck visible at <http://www.skendric.com/seminar>