

SharkFest '16

Troubleshooting a Multi-Tier Application in a Production Environment

June 16, 2016



Brad Palm
Lead Engineer | Marine Corps Network Efficiency Lab

DISCLAIMER

“The views expressed in this presentation are my own and do not necessarily represent the views of the United States Marine Corps, Department of Defense, or the United States Government.”

AGENDA

- Background
- Problem Framing
- Method
- Execution
- Analysis
- Questions

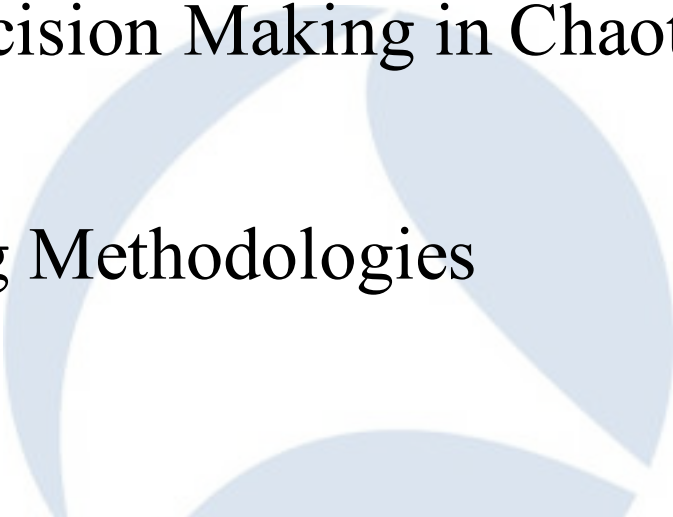




BACKGROUND

SharkFest '16 • Computer History Museum • June 13-16, 2016

BACKGROUND

- Training for Decision Making in Chaotic Environments
 - Troubleshooting Methodologies
 - Case Study
- 

Video of Chaotic Training at USMC Boot Camp - <https://www.youtube.com/watch?v=Z9-rhwHkP24>



SharkFest '16 • Computer History Museum • June 13-16, 2016

USMC Training for Chaos

- **BAMCIS**

- **B**egin planning, **A**rrange for Reconnaissance and Coordination, **M**ake Reconnaissance, **C**omplete Plan, **I**ssue Order, **S**upervise
- Six troop leading steps by which a leader plans and executes missions
- Logical and orderly process that makes the best use of time

- **Fog of War (Bits)**

- Uncertainty, hard to see patterns in midst of fog, separate signal from noise

- **Realistic Training**

USMC Training for Chaos



“Training to be comfortable in a range of environments and stressors”



Troubleshooting Methodologies

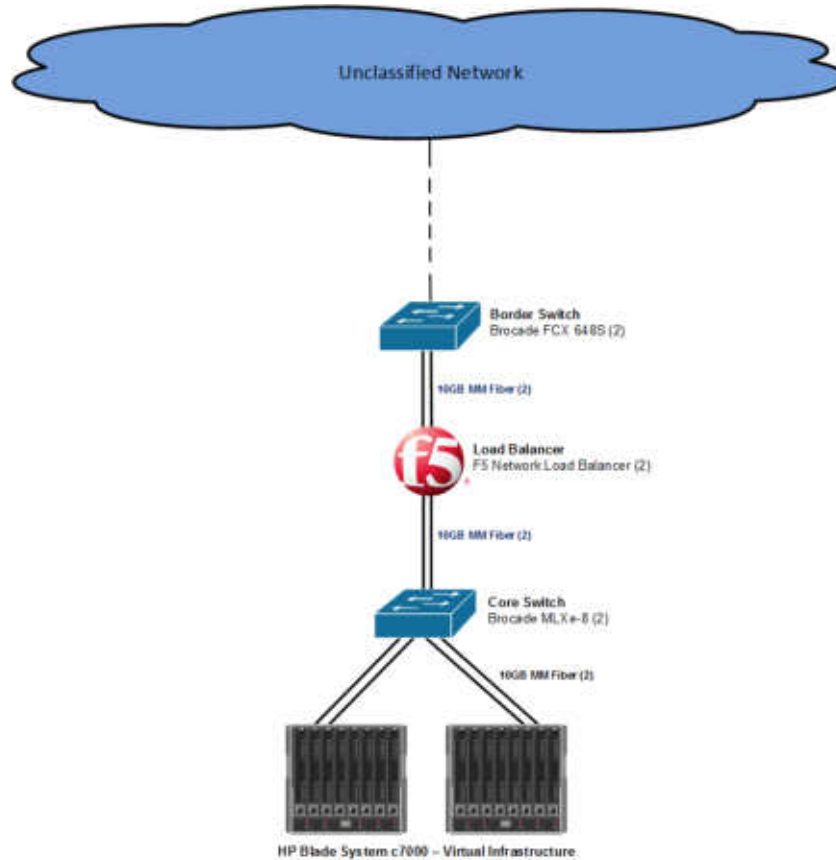
- Brainstorming
- 5 Whys
- Statistical Method
- Pattern Method
- Team Formation
- Cisco's Eight Step Method
- OSI Model



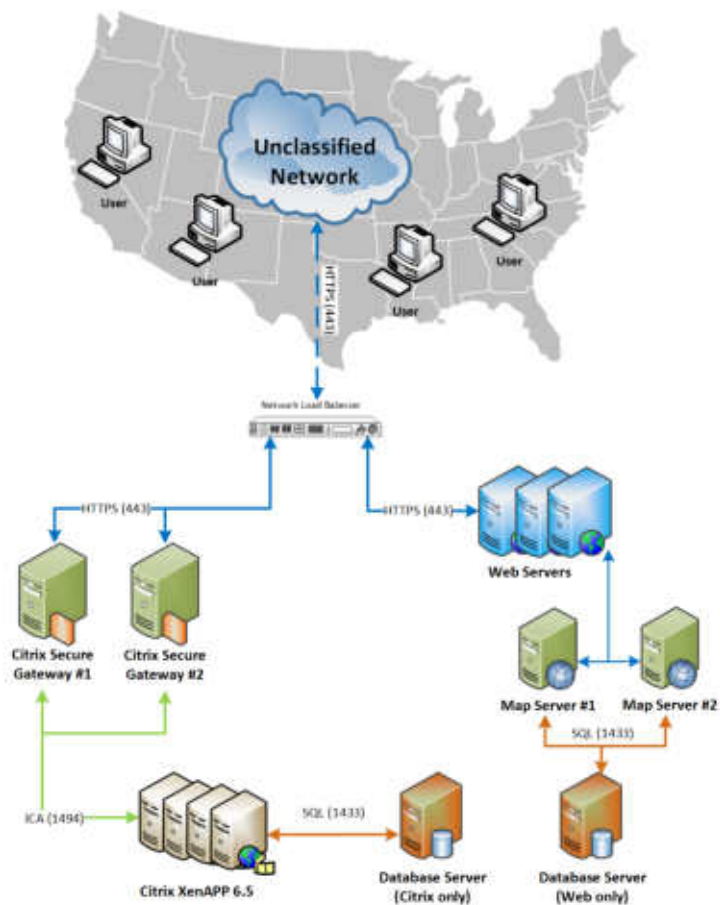
Case Study – 5Ws

- **Who** – Data center that hosts applications for MC wide consumption
- **What** – Users complaining about latency for an application that uses VDI to accomplish computationally intensive imagery manipulation
- **When** – Latency occurs at various times
- **Where** – Latency occurs for various users
- **Why** – ??? Help!!!

Case Study – Network Topology



Case Study – VM Architecture






PROBLEM FRAMING

SharkFest '16 • Computer History Museum • June 13-16, 2016

PROBLEM FRAMING

- How do we tackle complex problems and solve the WHY?
- Need solid foundation, that provides flexibility and facilitates multiple troubleshooting methodologies

Decision Making in Chaotic Environments



CAPTURE PLAN

Troubleshooting Methodologies

Problem Framing with Capture Plan

- Purpose
 - State objectives – define the problem
 - Objective requirements – what must be met to facilitate successful troubleshooting
 - Network metrics – what are we looking for based off probable cause
- Method
 - Tools
 - Capture location
- Execution
 - Team composition
 - TAP installation
- Analysis
 - Troubleshooting steps
 - Reporting



METHOD

SharkFest '16 • Computer History Museum • June 13-16, 2016

METHOD

- Purpose
 - State objectives
 - Objective requirements
 - Network metrics
- Method
 - Tools
 - Capture location
- Execution
 - Team composition
 - TAP installation
- Analysis
 - Troubleshooting steps
 - Reporting




Tools

- **TAPs** – Full Duplex, Aggregators, Copper, Optical, Conversion
 - Viavi nTAP, Profishark 1G, Ixia Net Optics
- **Capture Devices**
 - Netscout Network Time Machine (1 GB), Viavi Observer GigaStor (10 GB)
- **Analysis**
 - SteelCentral Packet Analyzer
 - Wireshark
- **Virtual Considerations** – VMs, Containers, Virtual TAPs


ixiaVIAYI
Fidelity NETWORK INSTRUMENTSNETSCOUT. | FLLIKE networks.riverbed

Capture Location

Network Topology



VM Architecture





EXECUTION

SharkFest '16 • Computer History Museum • June 13-16, 2016

EXECUTION

- Purpose
 - State objectives
 - Objective requirements
 - Network metrics
- Method
 - Tools
 - Capture location
- **Execution**
 - Team composition
 - TAP installation
- Analysis
 - Troubleshooting steps
 - Reporting



Team Composition

- Problem Manager
- Problem Solving Group
 - Mix of internal and external entities
 - Internal – network, virtual, DB, server
 - External – capture & analysis, Windows, VM
 - Need internal folks to be empowered, recognize when they don't have the authority, then be proficient in navigating the “policy” waters

TAP Installation

- In-Brief
- Common Concerns
- Installation
 - 300ms ~ 30 sec, Ethernet Autonegotiation
- Time Considerations
 - 24hr soak, during working hours/high probability of event



Demo

“Purpose of demo is to prove the following three items to instill confidence in the client that the TAP is not adverse to their network:

- Inserting TAP in a timely manner will not cause an outage to current TCP sessions*
- Fail to wire*
- Passive”*



“Demo showed a simple network between laptop and RaspberryPi, a persistent TCP session created through the use of a counter bash script and netcat, and an nTAP was inserted into the network to show the three items”



ANALYSIS

SharkFest '16 • Computer History Museum • June 13-16, 2016

ANALYSIS

- Purpose

- State objectives
- Objective requirements
- Network metrics

- Method

- Tools
- Capture location

- Execution

- Team composition
- TAP installation

- Analysis

- Troubleshooting steps
- Reporting



Troubleshooting Steps

1. **Define the Problem** – 5 Whys, Brainstorming, Problem Solving Group
2. **Gather Detailed Information** – Gain an understanding of architecture from in-house experts, supports Statistical and Pattern Methods
3. **Consider Probable Cause for Problem** – Discuss recent changes, investigate the network captures, identify the most likely hypotheses
4. **Devise Plan to Solve Problem** – Conduct preliminary/minimally disruptive testing, schedule testing time/location/equipment/personnel
5. **Implement the Plan** – Make necessary changes and capture test results
6. **Observe the Results of Implementation** – Confirm problem resolution, conduct regression testing
7. **Repeat Process if Plan Does Not Initially Solve Problem** – Re-define the problem, determine next hypothesis to test, *making changes incrementally*
8. **Document the Changes Made to Resolve the Problem** – Consolidate information gathered into network analysis report to communicate the process

Latency Analysis

- Troubleshooting latency:

- Network path
- Server side
- Client side



Statistical Method

Client 1 (Mid West)	WAN 1	LAN 2	VM 1	VM 2
Size	365.52 MB/ 564,158 pkts	334.62 MB/ 436,434 pkts	132.71 MB/ 290,417 pkts	2.1 GB/ 1,843,248 pkts
Average RTT	48 ms	47 ms	955 μ s	99 μ s
Duplicate ACK Errors	10,507 pkts	10,236 pkts	N/A	4,413 pkts
Retransmission Errors	4,172 pkts	4,030 pkts	N/A	744 pkts
Lost Segment Errors	160 pkts	154 pkts	61 pkts	140 pkts
Zero Window Errors	N/A	N/A	53 pkts	125 pkts
Reset Errors	6 pkts	3 pkts	N/A	17 pkts

Client 2 (Desert)	WAN 1	LAN 2	VM 1	VM 2
Size	26.03 MB/ 178,382 pkts	25.31 MB/ 178,326 pkts	130.96 MB/ 639,031 pkts	1.39 GB/ 1,872,794 pkts
Average RTT	70 ms	70 ms	700 μ s	168 μ s
Duplicate ACK Errors	3,769 pkts	3,769 pkts	3 pkts	3,364 pkts
Retransmission Errors	307 pkts	297 pkts	3 pkts	484 pkts
Lost Segment Errors	116 pkts	117 pkts	53 pkts	35 pkts
Zero Window Errors	29 pkts	29 pkts	640 pkts	49 pkts
Reset Errors	28 pkts	16 pkts	N/A	67 pkts

Statistical Method

Client 3 (East Coast)	WAN 1	LAN 2	VM 1	VM 2
Size	73.30 MB/ 294,128 pkts	72.13 MB/ 294,128 pkts	166.50 MB/ 490,632 pkts	2.65 GB/ 2,568,391 pkts
Average RTT	41 ms	41 ms	759 μ s	123 μ s
Duplicate ACK Errors	2,897 pkts	2,897 pkts	4 pkts	7,707 pkts
Retransmission Errors	792 pkts	792 pkts	5 pkts	1,289 pkts
Lost Segment Errors	245 pkts	245 pkts	160 pkts	266 pkts
Zero Window Errors	N/A	N/A	603 pkts	91 pkts
Reset Errors	64 pkts	64 pkts	N/A	35 pkts

Client 4 (East Coast)	WAN 1	LAN 2	VM 1	VM 2
Size	58.74 MB/ 232,503 pkts	57.72 MB/ 232,125 pkts	31.13 MB/ 135,806 pkts	254.97 MB/ 361,597 pkts
Average RTT	37 ms	37 ms	538 μ s	148 μ s
Duplicate ACK Errors	1,781 pkts	1,781 pkts	N/A	532 pkts
Retransmission Errors	549 pkts	545 pkts	1 pkt	66 pkts
Lost Segment Errors	171 pkts	171 pkts	84 pkts	43 pkts
Zero Window Errors	N/A	N/A	39 pkts	10 pkts
Reset Errors	33 pkts	25 pkts	N/A	25 pkts

Statistical Method

Client 5 (California/4G Commercial Network)	Client Machine
Size	4.22 MB/11,566 pkts
Average RTT	104 ms
Duplicate ACK Errors	736 pkts
Retransmission Errors	604 pkts
Lost Segment Errors	178 pkts
Out of Order Errors	139 pkts
Reset Errors	27 pkts
Missing Segments (ACKs)	2 pkts
Zero Window Errors	0 pkts

Network Path Latency Analysis

Client (Location)	Average RTT	User Experience	Notes
Client 1 (Mid West)	48 ms	Good	Physically located in the same building as the app, but traverses a network before hitting servers
Client 2 (Desert)	70 ms	Bad	Long time troubled area for app service.
Client 3 (East Coast)	41 ms	Bad	Long time troubled area for app service.
Client 4 (East Coast)	37 ms	Bad	Long time troubled area for app service.
Client 5 (West Coast)	104 ms	Good	Accessed app over a 4G hotspot with the highest RTT observed.

Network Path Latency Analysis



“What’s the big deal about 104ms, when we are used to 550ms on SATCOM!?!”




Client Side Latency Analysis

- Inability to capture client side
- Pattern Method:
 - GPO inspection
 - Standardized machine images, but many different ways of patching, updating, and policies taking effect
- Standardized Citrix GPO, agreed upon by all stakeholders

Server Side Latency Analysis

“Solution: TCP receive buffer at the secure gateway was the culprit and needed to be tuned. The zero window spikes coincided with application crashes for the users - these were confirmed during user tests”

No.	Stream Index	Time	TCP Stream Delta	rRTT	Packet Length	Source	Destination	Protocol	Info	Calculated window size
17472	7	33.056	0.00002000	1514	10.1.101.152	10.1.101.240	TCP	2598-27084 [ACK] Seq=820061 Ack=13892 Win=255 Len=1460	255	
17473	7	33.056	0.00003000	1514	10.1.101.152	10.1.101.240	TCP	2598-27084 [ACK] Seq=821521 Ack=13892 Win=255 Len=1460	255	
17474	7	33.056	0.00002000	1514	10.1.101.152	10.1.101.240	TCP	2598-27084 [ACK] Seq=822981 Ack=13892 Win=255 Len=1460	255	
17475	7	33.056	0.000062000	60	10.1.101.240	10.1.101.152	TCP	27084-2598 [ACK] Seq=13892 Ack=824441 Win=6 Len=0	6	
17476	7	33.056	0.000091000	1514	10.1.101.152	10.1.101.240	TCP	2598-27084 [ACK] Seq=824441 Ack=13892 Win=255 Len=1460	255	
17480	7	0.007	0.006134000	66	10.1.101.240	10.1.101.152	TCP	[TCP ZeroWindow] 27084-2598 [PSH, ACK] Seq=13892 Ack=8259	0	
17486	7	0.007	0.007793000	66	10.1.101.240	10.1.101.152	TCP	[TCP ZeroWindow] 27084-2598 [PSH, ACK] Seq=13904 Ack=8259	0	
17488	7	0.007	0.000115000	60	10.1.101.152	10.1.101.240	TCP	2598-27084 [ACK] Seq=825901 Ack=13916 Win=255 Len=0	255	
17492	7	0.015	0.007410000	60	10.1.101.240	10.1.101.152	TCP	[TCP Window Update] 27084-2598 [ACK] Seq=13916 Ack=825901	256	





SharkFest '16 • Computer History Museum • June 13-16, 2016