



Tivoli Software

# Identifying and Solving Network Performance Problems on zEnterprise



SHARE Boston  
August 4, 2010

Mac Holloway (mhollowa@us.ibm.com)

Dean Butler (butlerde@us.ibm.com)

© 2010 IBM Corporation



# Agenda

- Intro
- Problem spaces
- Configuration issues
- Operational issues
- Modern Applications
- Scenarios
- Summary



# Intro

- Host network management definition
  - Enabling effective use of z/OS Communications Server, VTAM and OSA communications as part of End-To-End solutions in IT shops.
- Presenters
  - **Dean Butler** – 15+ years of network and network management development both inside and outside IBM. Currently a System z software architect in Tivoli.
  - **Mac Holloway** – 20 years of IBM networking in NSD, NHD and Tivoli including work on NWAYS, zNV, Mainframe Networks and z/VM and Linux



Tivoli Software

# Mainframe Networking Performance Problem spaces





# Problem spaces

## ■ What we hear

- “A critical application is “broken”. We all get on a bridge call/line. Everyone says “Mine stuff is okay. It must be the network. I need to be able to say it is not the network or at least not my part of the network.”

## ■ Configuration issues

- This is the most common type of problem we see. This is the most common type of problem zCS sees
- Examples of VTAM, zCS, SNMP, IPSec configuration issues
- An approach

## ■ Operational issues

- Problem sources
  - Protocol – malformed packets, unusable ports, ...
  - Resources – packets dropped, buffer overflows
  - Indirect – packet reassembly, response time, rate changes
- Location indicators
  - Outboard – from the MAC out into the network and beyond
  - Stack – from the MAC to the buffer interfaces above TCP and UDP
  - App – above the stack including FTP, TN3270



Tivoli Software

# Mainframe Networking Performance Configuration Issues





## Configuration Issues - examples

- OSA
  - OSA-Express Direct SNMP subagent (IOBSNMP) or OSA/SF application (IOAOSASF) and the OSA/SF sockets application (IOASNMP) – running ?
  - RACF for OSA/SF SNMP sub-agent (IOASNMP) – security messages
- IPSec
  - IKE Daemon started ?
  - PAGENT Daemon started ?
  - IPSec NMI access authorized ?
- TN3270 & FTP
  - z/OS Communications Server real-time SMF data network NMI enabled ?
  - Monitoring app authorized ?
  - Sliding window or bucket count data – configured in Telnet profile ?
- SNA NMI
  - zCS SNA NMI not enabled, ?
  - OMVS segment created for VTAM ?
  - Monitoring app authorized



## Configuration Issues - examples

### ■ VTAM

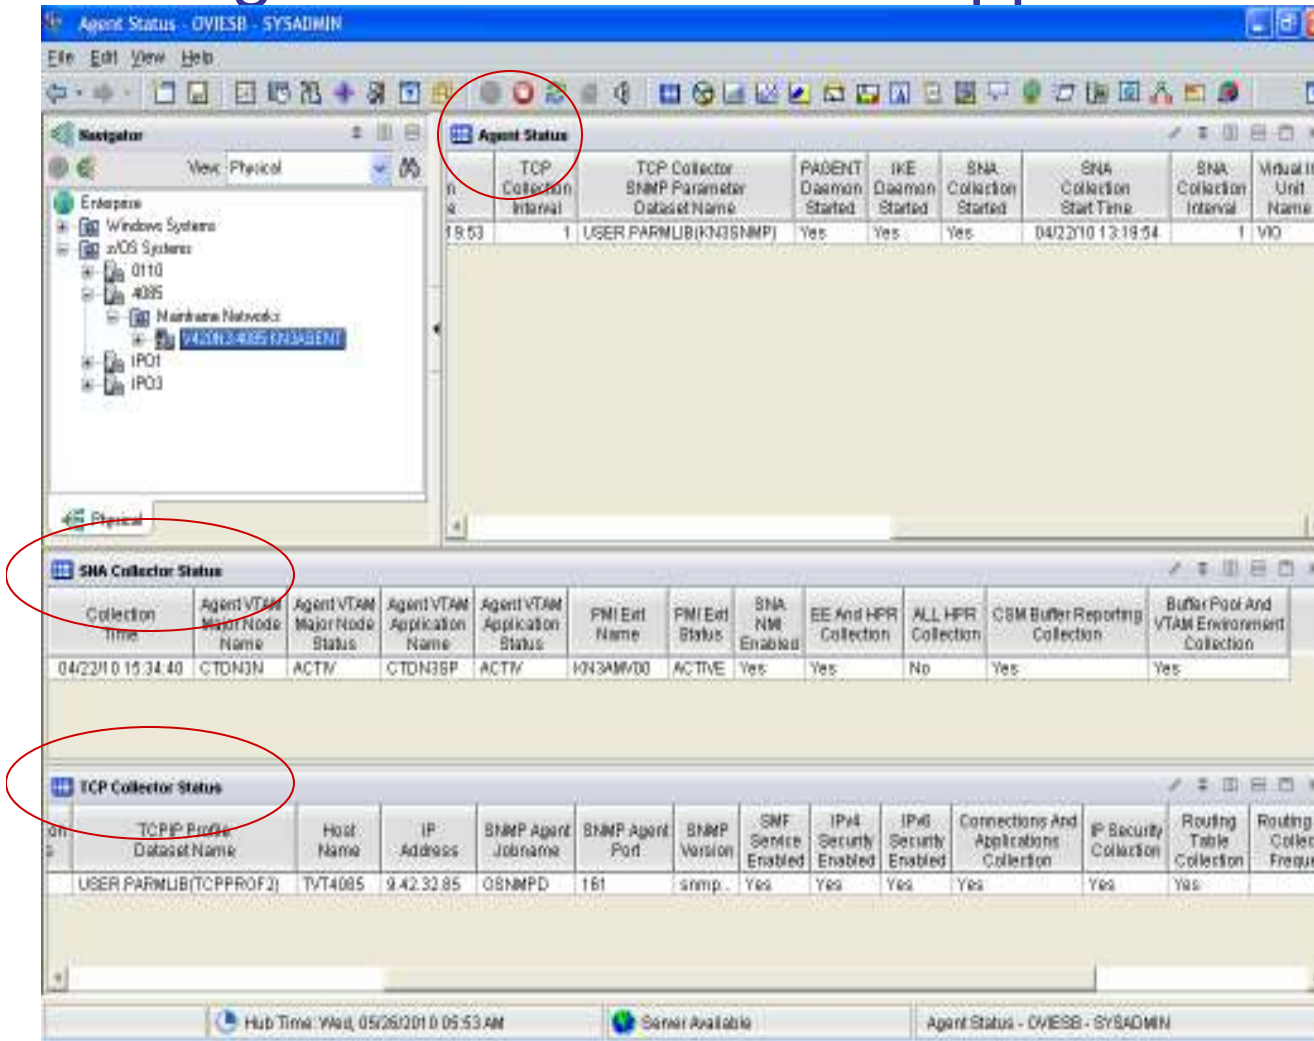
- Is it running ? Have you done a vary to activate it ?
- Is SNA data collection configured ?
- Is your monitoring app in the VTAMLST ?
- Is the PMI exit available to VTAM ? Add DD card, quiesce VTAM, restart VTAM

### ■ SNMP

- Is it configured ? SNMPD.CONF
- Is it running ? OSNMP
- Are you accessing the right address ? loopback address (127.0.0.1)
- Do you have the right community name ? Check SNMPD.CONF
- Are you using the right port ? 161
- SNMP requests are timing out ? Your application



# Configuration Issues – an approach



**Agent Status**

TCPIP Profile	TCP Collector	TCP Collector	PAGENT	IKED	SNA	SNA	SNA	Virtual ID
Dataset Name	Collection Interval	SNMP Parameter Dataset Name	Daemon Started	Daemon Started	Collection Started	Collection Start Time	Collection Interval	Unit Name
USER.PARMLIB(TCPPROF2)	18.53	1	USER.PARMLIB(KN3SNMP)	Yes	Yes	Yes	04/22/01 13:19:54	1 VIO

**SNA Collector Status**

Collection Time	Agent VTAM Major Node Name	Agent VTAM Major Node Status	Agent VTAM Application Name	Agent VTAM Application Status	PMI Ext Name	PMI Ext Status	SNA NMI Enabled	EE And HPR Collection	ALL HPR Collection	CSM Buffer Reporting Collection	Buffer Pool And VTAM Environment Collection
04/22/01 15:34:40	CTDN3N	ACTV	CTDN3BP	ACTV	KN3AMV00	ACTIVE	Yes	Yes	No	Yes	Yes

**TCP Collector Status**

TCPIP Profile Dataset Name	Host Name	IP Address	SNMP Agent Jobname	SNMP Agent Port	SNMP Version	SNMP Service Enabled	IP4 Security Enabled	IP6 Security Enabled	Connections And Applications Collection	IP Security Collection	Routing Table Collection	Routing Table Collection Frequency
USER.PARMLIB(TCPPROF2)	TVT4085	9.42.32.85	OSNMPPD	161	snmp	Yes	Yes	Yes	Yes	Yes	Yes	

Hub Time: Wed, 05/26/2001 05:53 AM  
 Server Available  
 Agent Status - OVIESB - SYSADMIN

Status on:

- SNA NMI
- VTAM
- PMI
- PAGENTD
- IKED
- OSNMP
- ...

Situations/actions based on status



Tivoli Software

# Mainframe Networking Performance Operational Issues





# Operational Issues Indicators


## Types of indicators

- Protocol errors – malformed packet, unmonitored port
- Resource constraints – dropped packets, buffer overflows
- Indirect – packet reassembly, response time. rates

## Direction of indicator

- Outboard – from the MAC out into the network and beyond
- Stack – from the MAC to the buffer interfaces above TCP and UDP
- App – above the stack including FTP, TN3270


# The layers – a quick look




## Operational issues – more detail



Each of these parts has data that can indicate



## Protocol indicators



### TCP/IP

Input Discards  
Output Discards  
UDP Discard  
UDP Input Errors  
UDP No Port

### OSA


Fragments, jabber, length error, CRC, alignment  
Unknown IP Frames

### Interfaces

Inbound Packet – discarded, in error  
Outbound Packet – discarded, in error  
Utilization  
Transmission Rates  
Unknown IP Frames



## Resource constraint indicators



### **TCP/IP**

CPU Percentage

CSA – Allocated, In Use

Authorized Private Storage – Allocated, In Use

ECSA storage – max, allocated, in use, pools

Datagrams Discarded

Backlog Connections Rejected

### **UDP**

Datagrams Discarded

### **HPR**

Throughput rate – allowed

Unacknowledged Buffers – high water mark

### **Interfaces**

Receive/Transmit Bandwidth utilization

### **OSA**

PCI Utilization

Processor Utilization

Missed Packets


### **VTAM**

CPU Percentage

CSA – Allocated, Allowed, In Use



## Indirect indicators



### TCP

Percent Segments Retransmitted  
Response Time  
Segments Retransmitted  
Fragmentation  
Reassembly  
TCP Retransmit  
Out of Order Segments  
Segments Retransmitted  
Remote Window Size Frequency  
Response Time Variance  
TCP Keep-Alive Drops

### HPR

Out of Sequence Buffers  
Packet Retransmission Rate  
Path Switches  
Response Time Variance  
Smoothed Round Trip Time

### TN3270

Average IP Response Time and variance  
Average SNA Response Time and variance







Tivoli Software

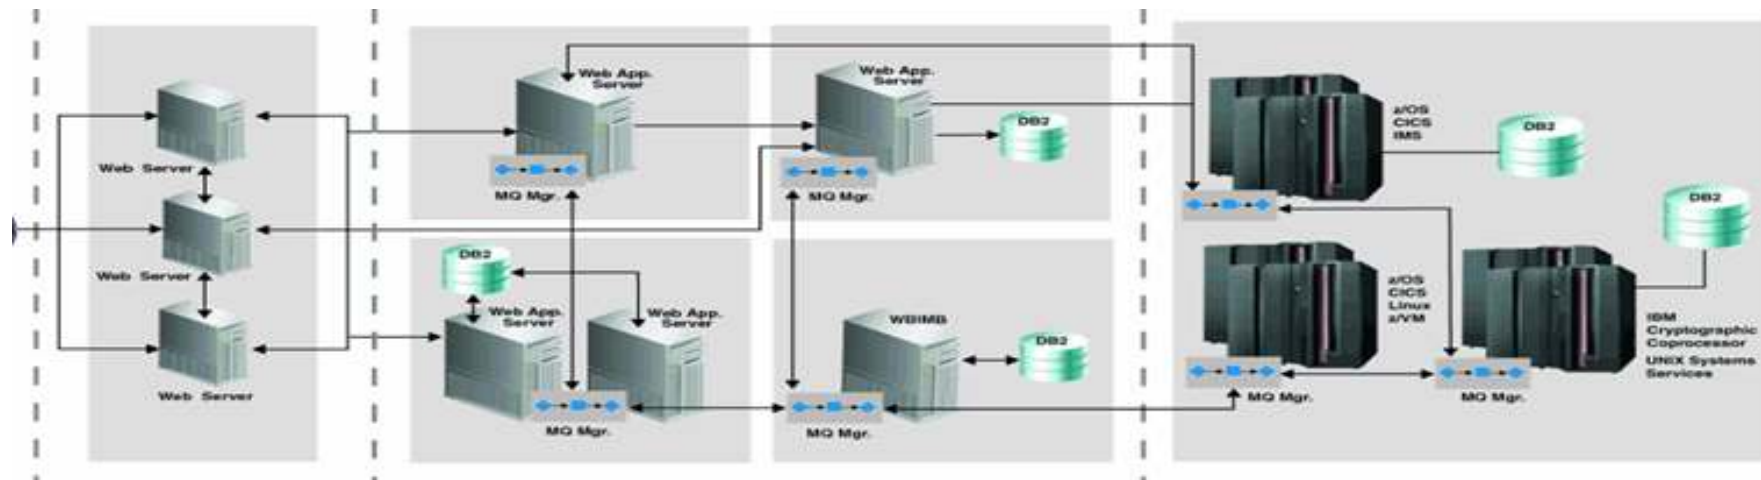
# Modern Applications ... Integrated Management



# Workflow for Resolving Composite Application Problems



**SHARE**  
Technology • Connections • Results



**Sense**

Detect that a threshold has been breached and that a problem occurred, or is about to happen

**Isolate**

Pinpoint the problem to a specific part of the environment and hand-off to the appropriate specialist


**Diagnose**

Drill down into the details and get to the root cause of the problem

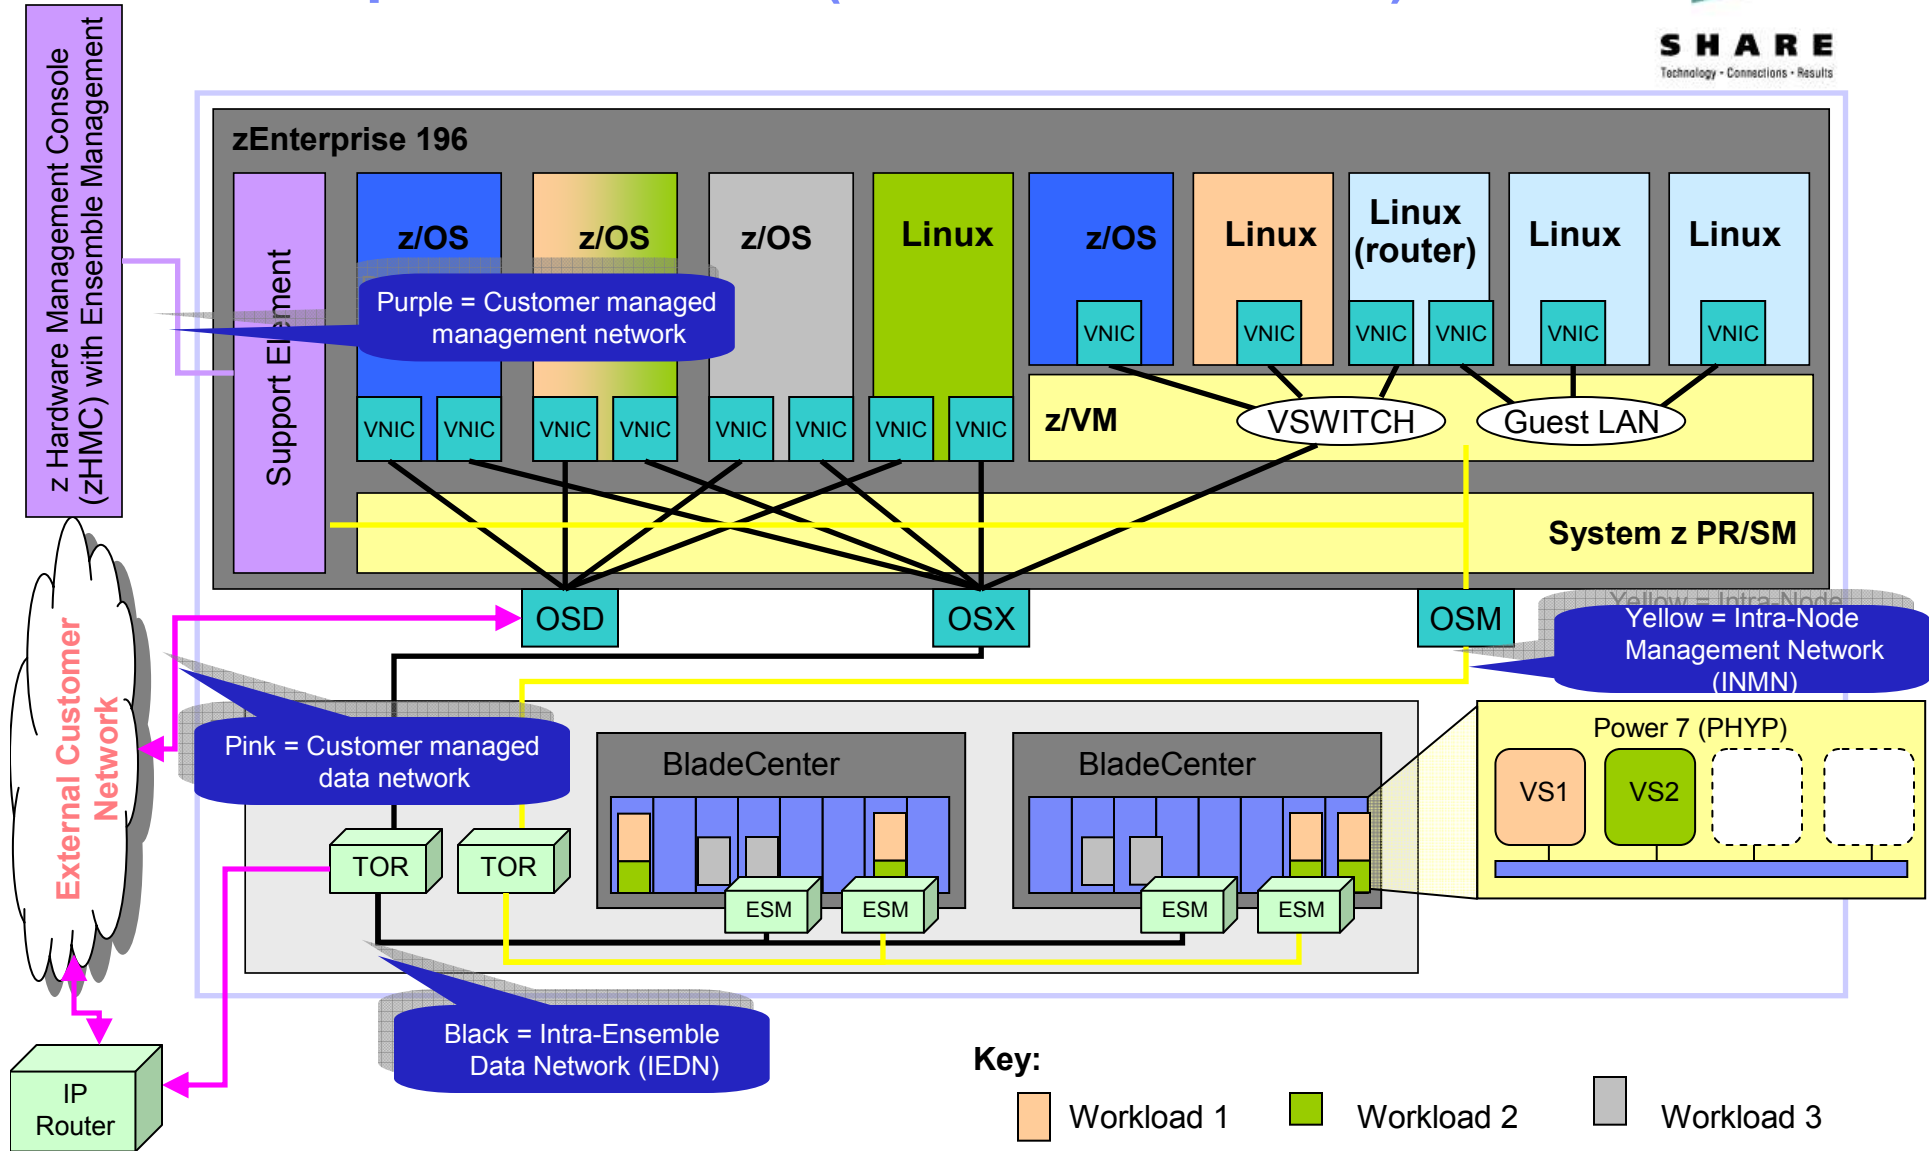
**Repair**

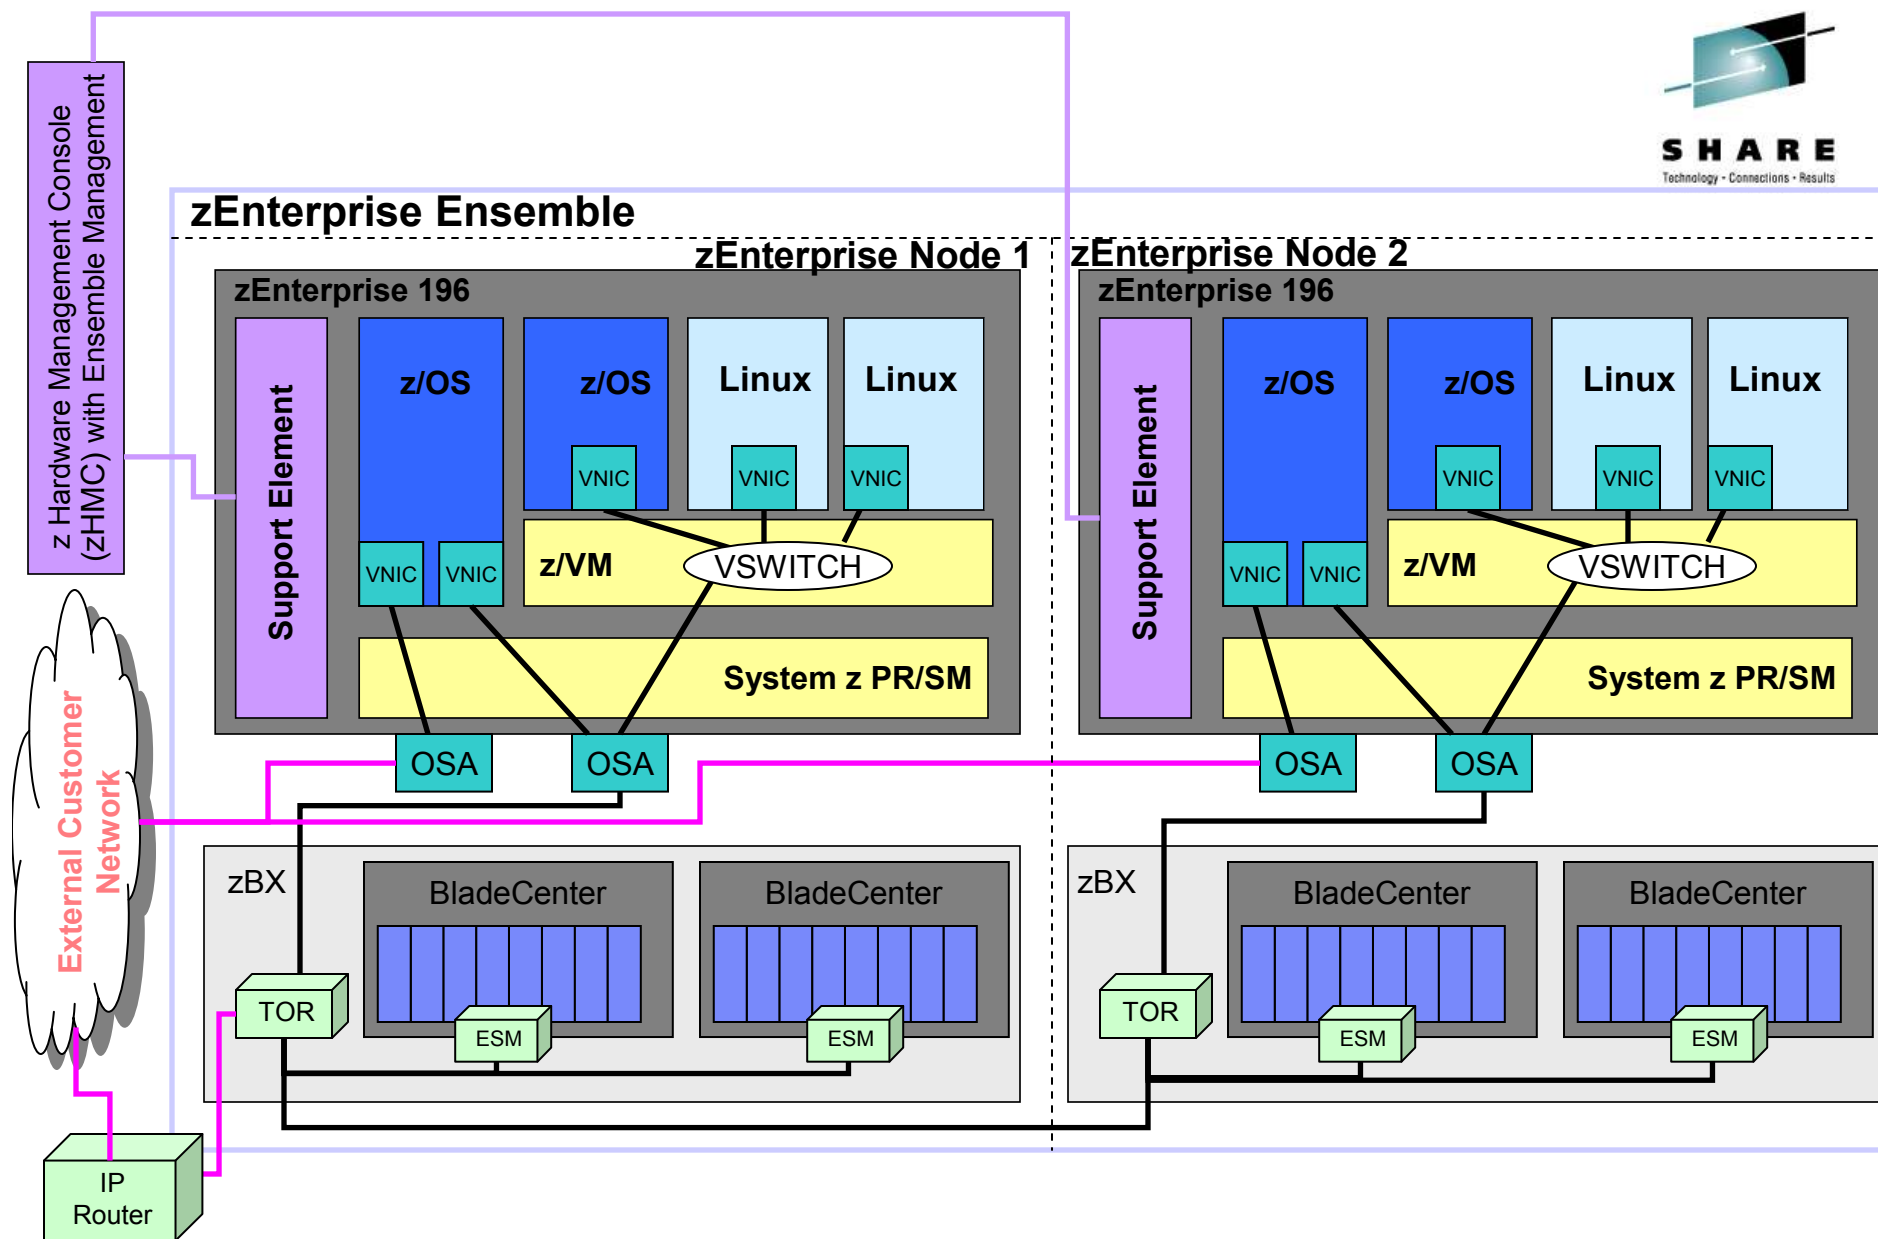
Fix the faulty component, validate the fix and roll back into production

# zEnterprise with zBX (z Blade Extension)

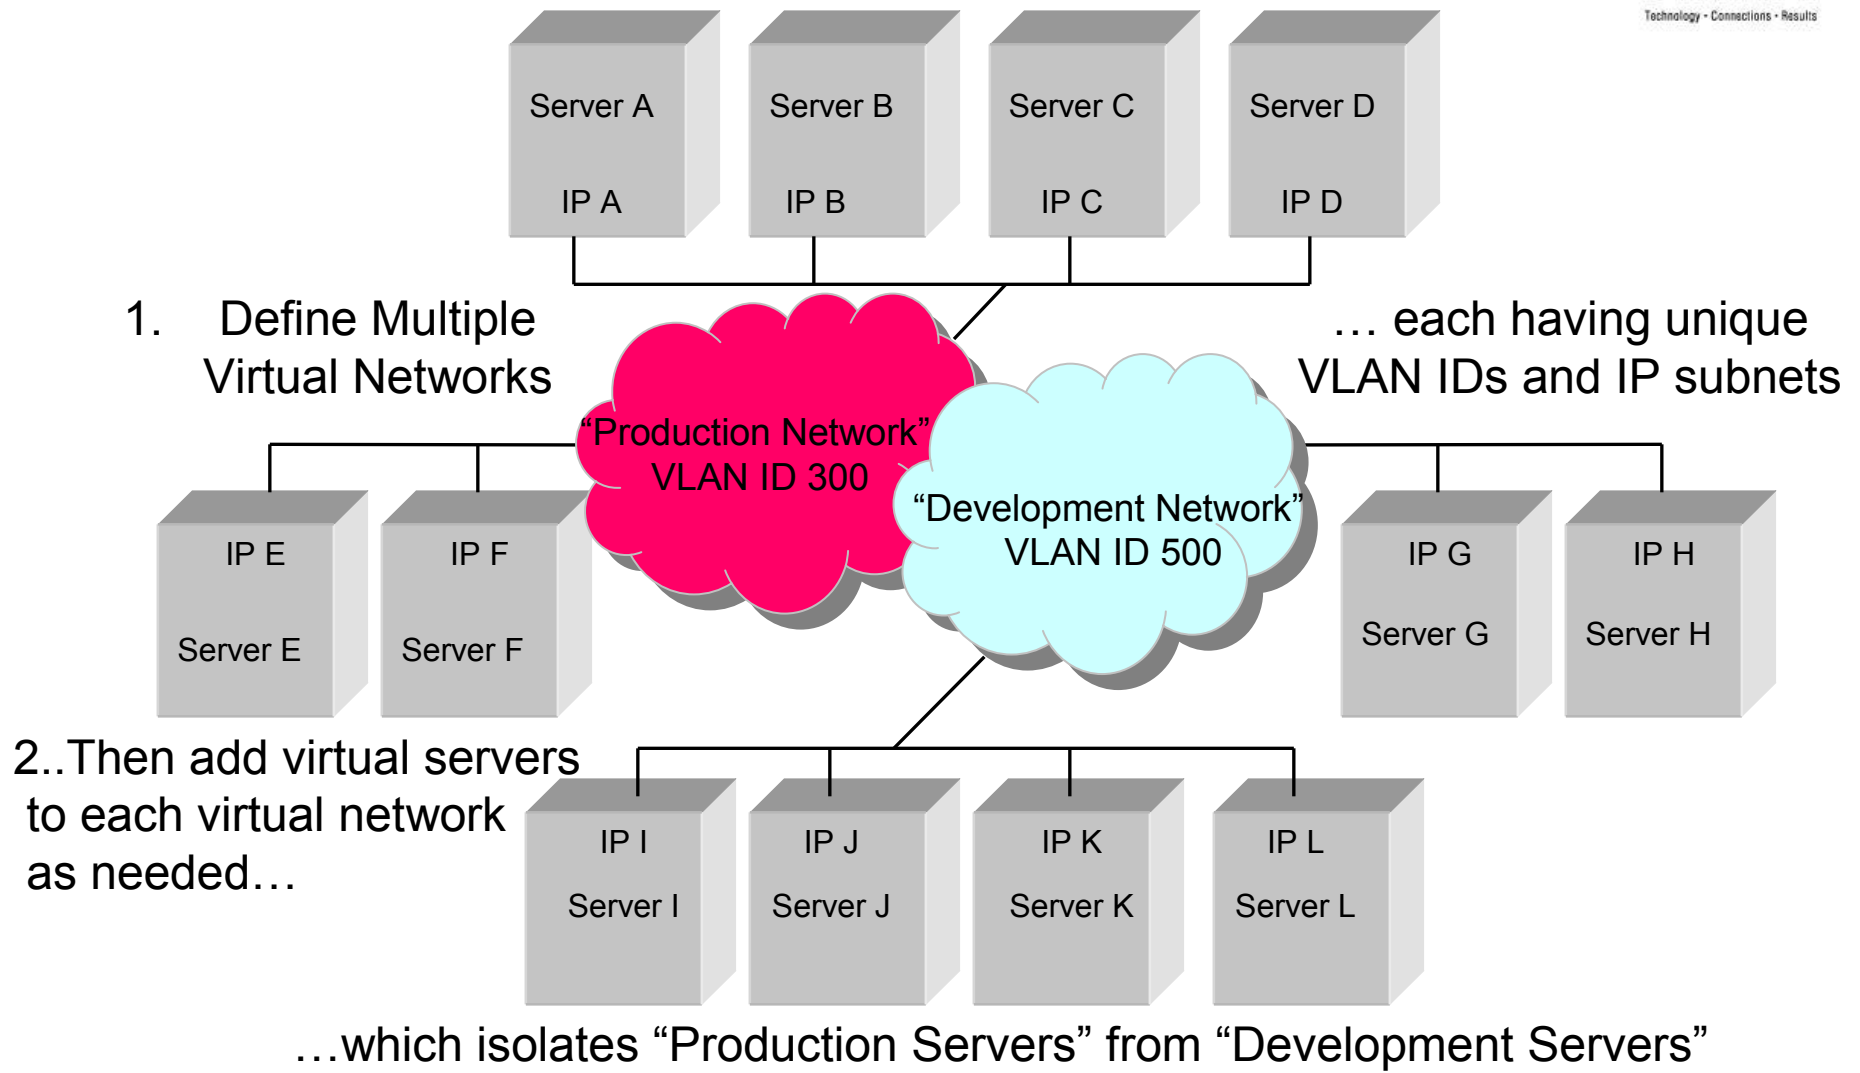


**SHARE**  
Technology • Connections • Results



SHARE  
Technology • Connections • Results**zEnterprise Ensemble****zEnterprise Node 1****zEnterprise Node 2**

# Multiple Virtual Networks – Isolation



## Which zEnterprise network components are of interest from a monitoring perspective?



- OSAs – Support is already available to monitor OSAs today.
  - New OSA types (OSX and OSM)
- A private and physically isolated management network (the intranode management network - INMN), connects all zEnterprise System resources (CPCs, BladeCenters, etc.) for management purposes. This INMN is pre-wired, internally switched, configured, and managed with full redundancy for high availability.
  - Throughput, dropped packets, fragmentation, etc are NOT of interest
- A private and secure OSA-Express Ethernet intraensemble data network (IEDN) that connects all elements of a zEnterprise System ensemble. The IEDN is access-controlled using integrated virtual local area network (VLAN) provisioning. IEDN management provides enforcement of strict access control across heterogeneous environments, further augmenting security and simplicity.
  - Throughput, dropped packets, fragmentation, etc are of interest
- VLANs – Defined in zHMC. Includes one or more virtual servers per VLAN ID. Key performance metrics (throughput, dropped packets, fragmentation, etc) would be of interest per interface.



Tivoli Software

# Scenarios








## Scenario A: Its not the Network!

### The setting:


A company relies on batch FTP to copy files between a mainframe at headquarters and each of its retail stores every night (local store time). Sales and inventory data is uploaded and product and pricing changes are downloaded to the stores. One morning, a systems administrator notices that some of the FTP jobs have not finished. He reports the problem to the IT help desk. The problem is routed to the mainframe networks systems programmer.





## User reports batch FTP failures




- Start with checking current activity: FTP transfers & FTP sessions

	Collection Time	Application Name	FTP Type	Remote IP Address	Remote Port	Local IP Address	Local Port	User ID on Server	Client User ID	Session Start	Session End	Session Duration
	08/11/08 23:52:22	FTPD1	Server	9.65.126.164	3000	9.42.45.179	21	USER2		08/11/08 23:52:18		0

	Collection Time	Remote IP Address	Remote IP Port	Local IP Address	Local IP Port	User ID on Server	Client User ID	Role	Transmission Start	Transmission End	Transmission Duration	Bytes Transmitted (in GB)	Bytes Transmitted
	08/11/08 23:53:03	9.65.126.164	3006	9.42.45.179	20	USER2		Server	08/11/08 23:52:49	08/11/08 23:53:01	11420	0	1440054
	08/11/08 23:53:39	9.65.126.164	3010	9.42.45.179	20	USER2		Server	08/11/08 23:53:39		0	0	0

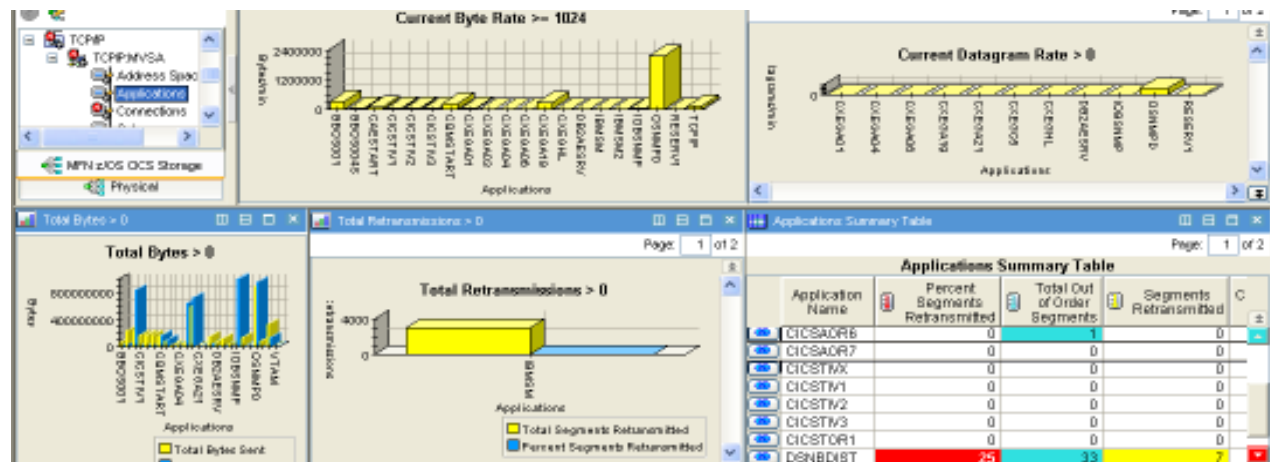


**SHARE**  
Technology • Connections • Results

# Check Applications and Connections

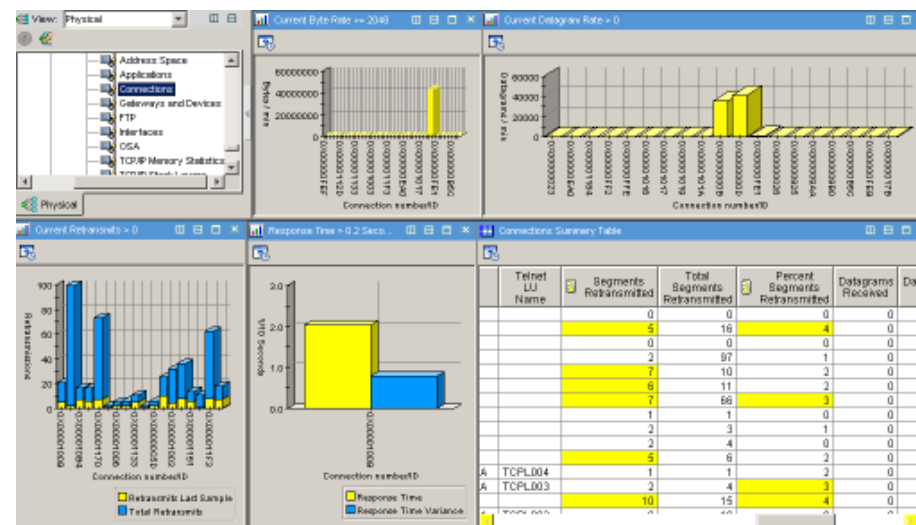
## Applications:

- Accepting connections?
  - Rate, Backlog, Rejections
- Last activity time
- Response Times
- Retransmissions
- Transmit / Receive Rates
- Out of order segments
- CICS, IMS, WAS, z/OS



## Connections:

- Start time/duration
- Response Time
- Response Time Variance
- Retransmissions
- Transmit / Receive Rates
- Out of order segments

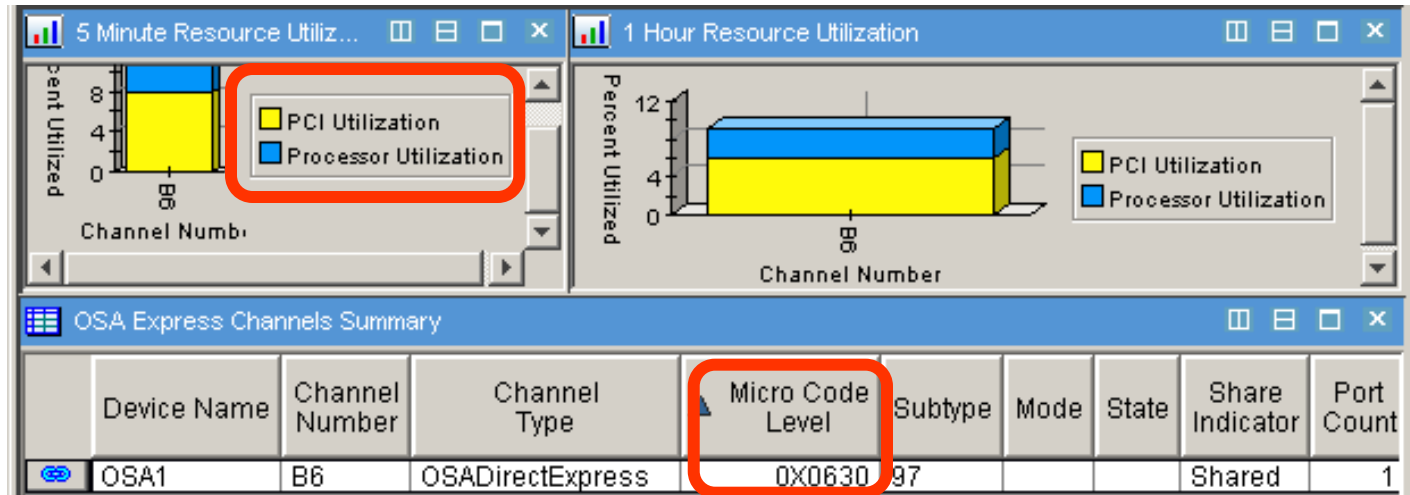




## Check OSA and Interfaces

### OSA

- Online Status
- Configuration
- Microcode Level
- Utilization
- Transmission Rates
- Unknown IP Frames
- By LPARS
- By Ports



### Interfaces

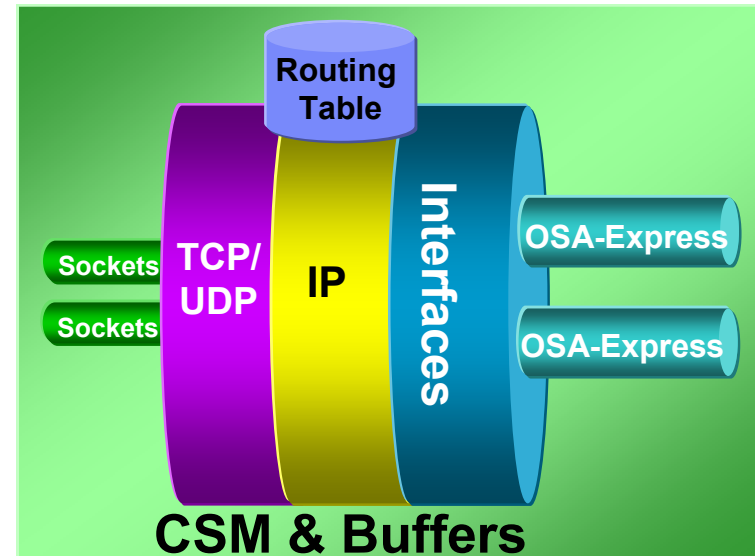
- Packet Errors
- Bandwidth Utilization
- MTU Size

	Interface Name	Description	Interface Type	Current State	MTU Size	Transmit Packet Rate	Receive Packet Rate
🌐	TCPIP LINK	IP Assist QDIO Ethernet	ethernetCsmacd	Up	1492	4312	74909
🌐	LOOPBACK	Loopback	softwareLoopback	Up	65535	890	890
🌐	LOOPBACK	Loopback Device	propVirtual	Up	0	890	890
🌐	OSA1	Multipath Channel IP Assist Device	propVirtual	Up	0	4312	74909
🌐	F7AXCFSA	Multipath Channel Print-to-Print	mnc	Down	55296	0	0



## Check TCP/IP Stack

- Retransmits - Network congestion
- Out of Order - Routing or network congestion
- Fragmentation - MTU size
- Discards - Resource constraints
- Timeouts - Connectivity
- UDP input Errors - Attack
- UDP Discards - Wrong Sockets
- High storage utilization - Could indicate Network congestion
  - This can result in requests backing up in storage

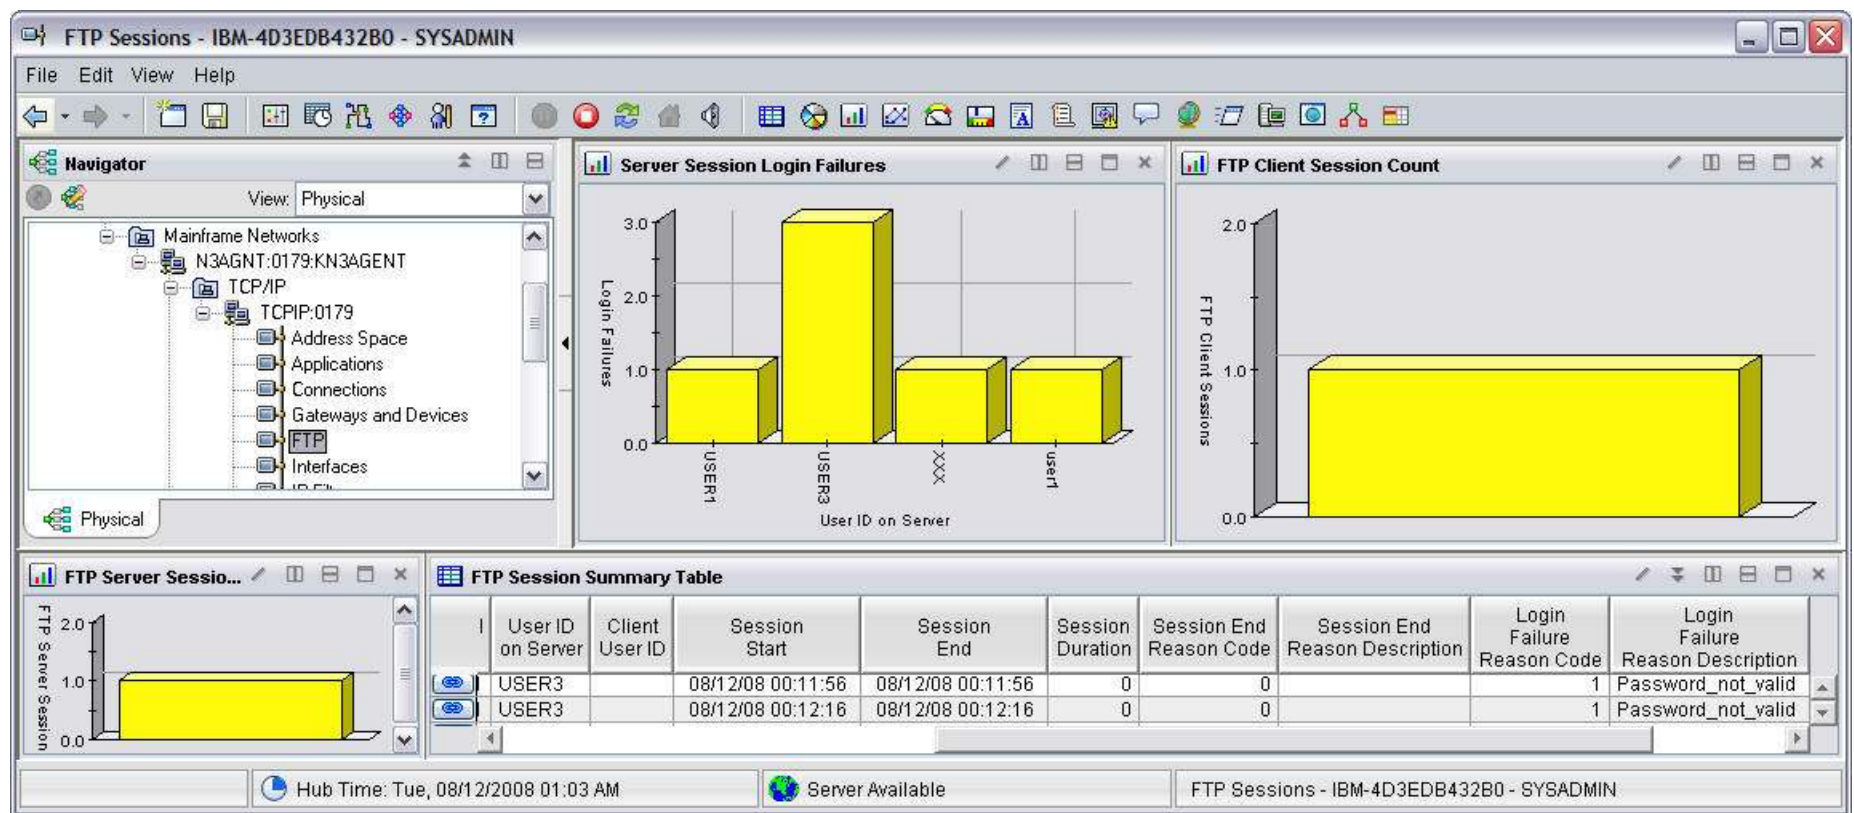


Output Routes	Reassembly Count	Reassembly Percentage	Reassembly Failure Count	Reassembly Failure Percentage	Fragmentation Count	Fragments To Be Reassembled	Fragmentation Percentage
0	6453	0	256	2	0	0	



## Network is OK, then where is the problem?

Check job logs, SYSLOG, and syslogd for clues.  
Dataset access issue? FTP session login failure?








## Scenario B: Slow response time in web service

### The setting:

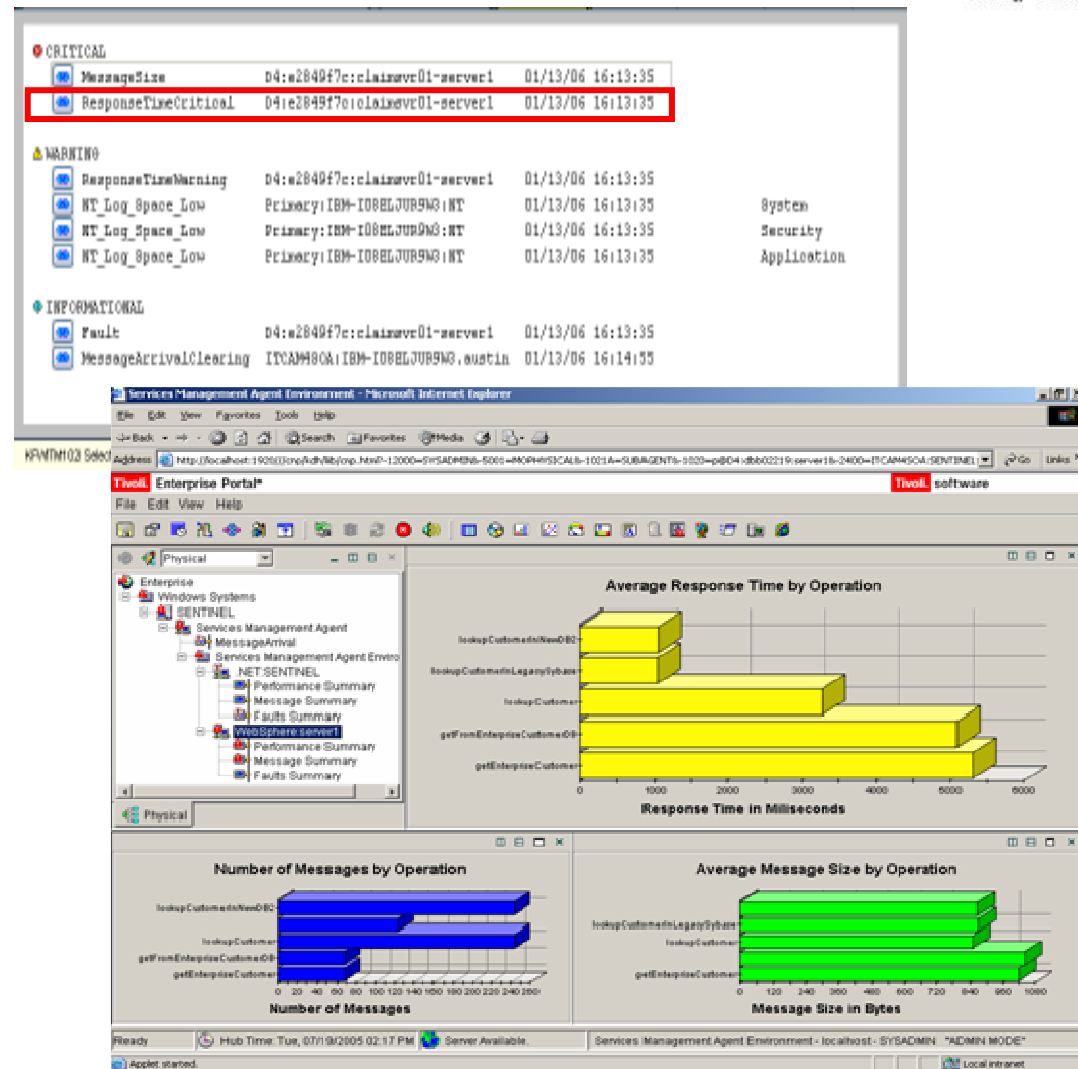
A company recently deployed a set of web services that replaced a very high profile application. The operations team monitors the performance closely. When performance degrades, its time to investigate...



**SHARE**  
Technology • Connections • Results

## Slow response time in web service ...

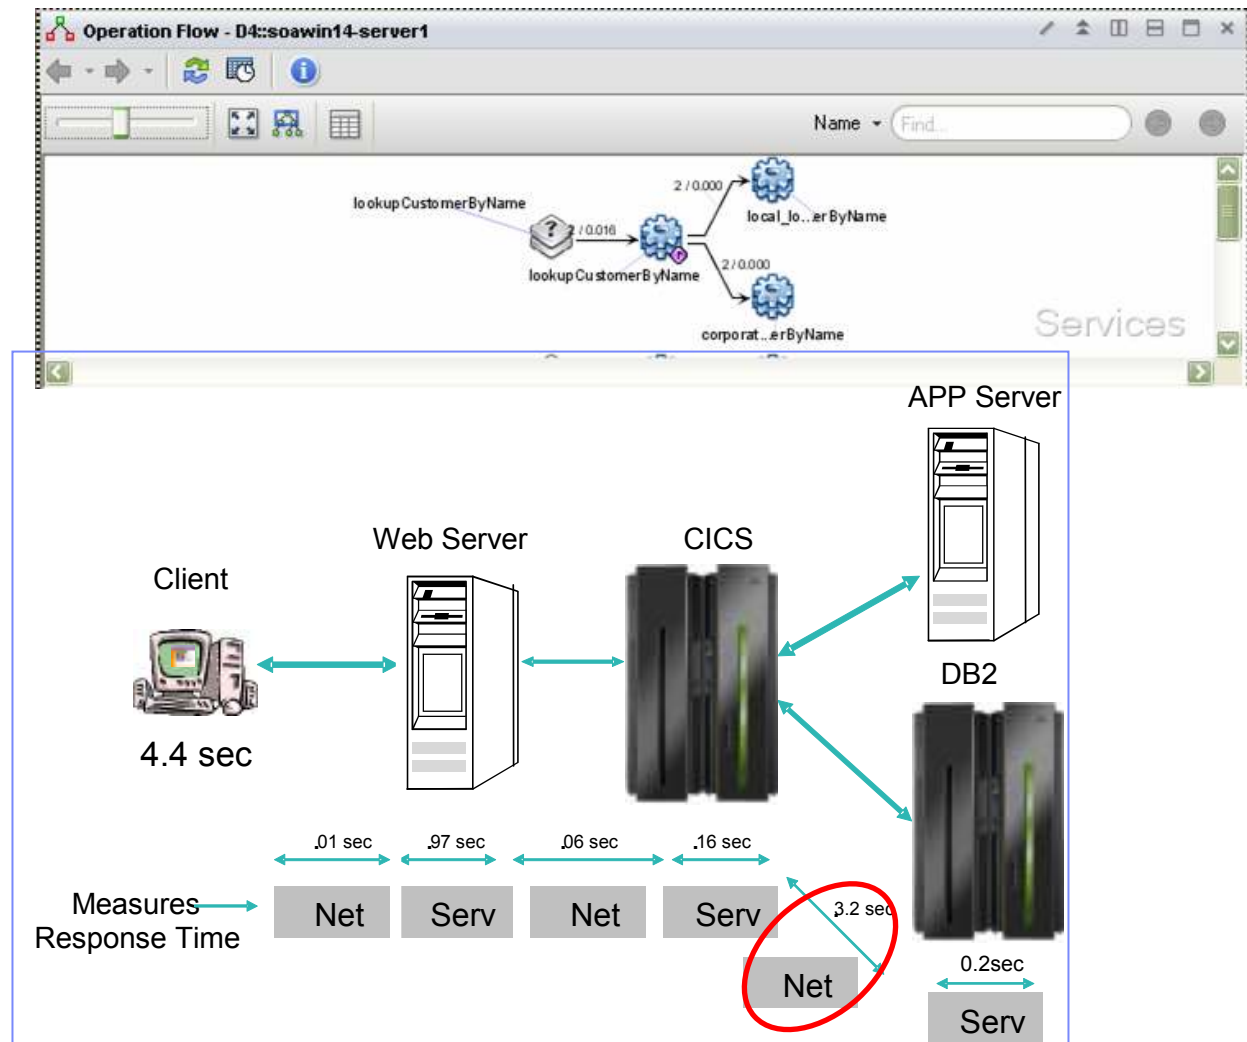
1. An alert identifies a response time problem. Annette, an operator, determines that slow response times are being recorded for the new web services.
2. Annette checks the number of requests and the message size and determines this is a normal volume of traffic. Annette passes the issue to Johann, a SME.





## Slow response time in web service ...

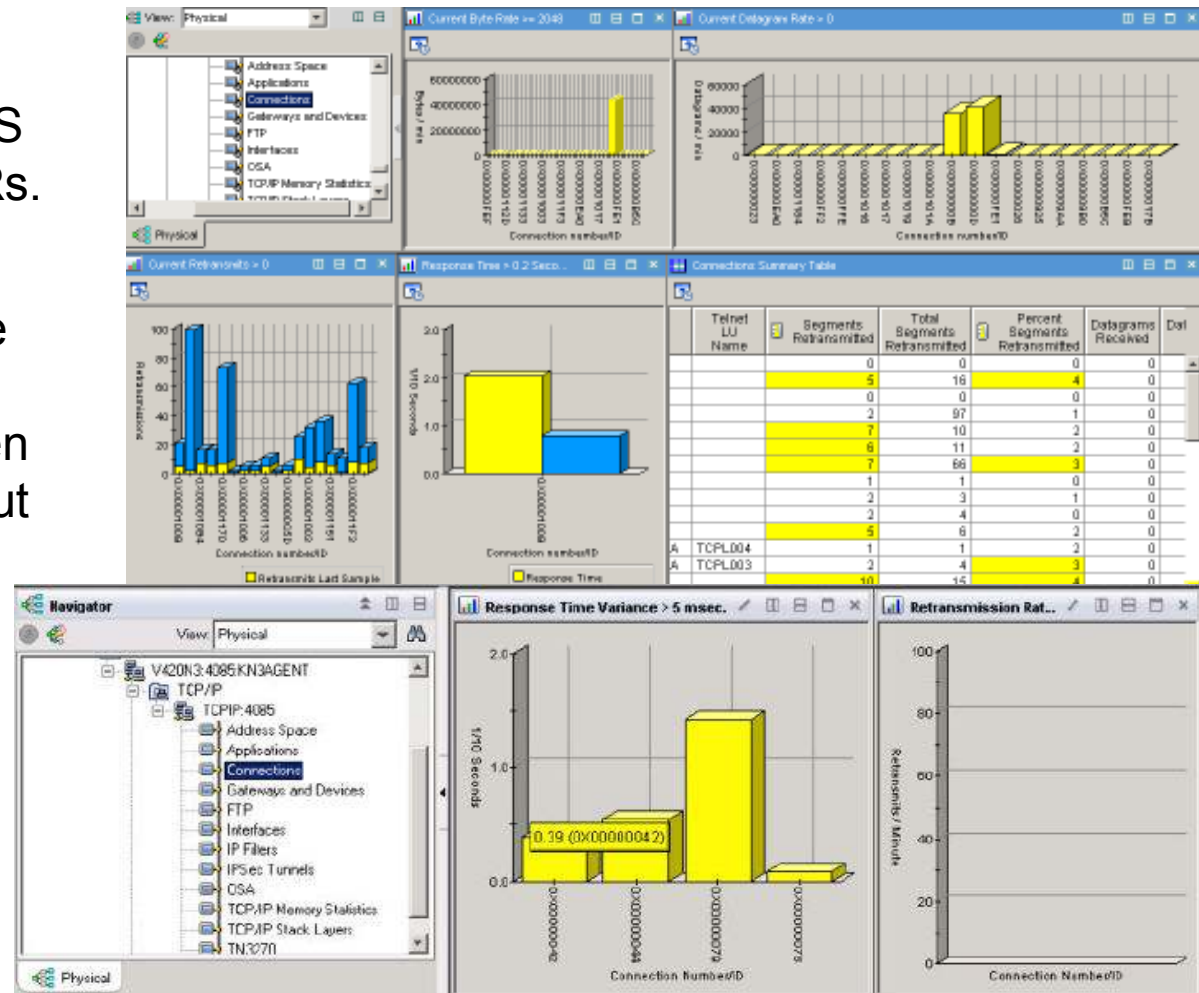
3. Johann begins by looking closer at the web services. Identifies flows and response time for each step.
4. Problem appears to be with the network between the CICS and DB2 servers. These two LPARs are connected by a data center network.





# Slow response time in web service ...

5. Johann views metrics for connections between CICS and DB2 on the two LPARs.
6. Johann notices there have been retransmits and out-of-order segments between CICS and DB2 servers. But what is the root cause?

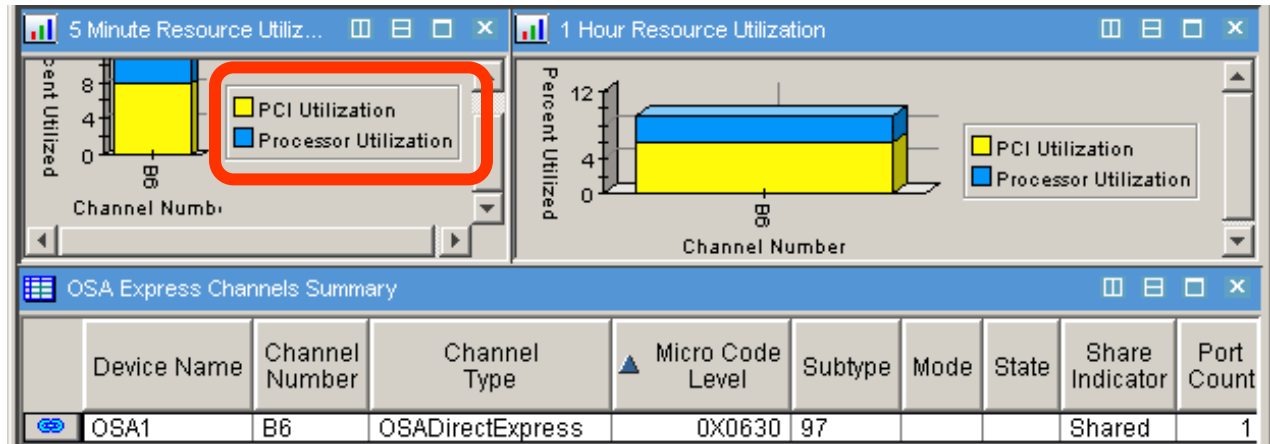




## Slow response time in web service ...

7. Johann checks the OSA cards and discovers the OSA on the DB2 server has high PCI and processor utilization.
8. Further checks reveal contention on OSA with other LPARs in the CEC is causing the performance issues.

Each OSA is dedicated to an LPAR, but also serves as backup OSA for a 2<sup>nd</sup> LPAR. Switch other LPAR to its primary OSA.



OSA-Express Ports Summary								
Collection Time	Port Name	Link Name	Port Type	Hardware State	Disabled Status	Service Mode	Configuration Name	Configuration Speed
12/16/08 13:34:34	OSAA	TCPIP.LIN	oneThousandBaseTEthernet	enabled	0X0000	NotInServiceMode	IBM Default ConfigFile: 1000Base T	oneThousandMbFullDuplex

LPAR Usage Summary for Channel B6											
Collection Time	LPAR Name	LPAR Logical Channel Subsystem	LPAR Number	LPAR Status	Processor Utilization Per Minute	Kilobyte Rate In Per Minute	Kilobyte Rate Out Per Minute	Processor Utilization Per Five Minutes	Kilobyte Rate In Per Five Minutes	Kilobyte Rate Out Per Five Minutes	Processor Utilization Per Hour
12/16/08 13:14:27			0	1 unknown	0	0	0	0	0	0	0
12/16/08 13:14:27			1	1 unknown	0	0	0	0	0	0	0
12/16/08 13:14:27			0	2 unknown	0	0	0	0	0	0	0
12/16/08 13:14:27			1	2 unknown	0	0	0	0	0	0	0
12/16/08 13:14:27			0	3 unknown	0	0	0	0	0	0	0
12/16/08 13:14:27			1	3 unknown	0	0	0	0	0	0	0
12/16/08 13:14:27			0	4 unknown	0	0	0	0	0	0	0
12/16/08 13:14:27			1	4 unknown	0	0	0	0	0	0	0
12/16/08 13:14:27			0	5 unknown	0	0	0	0	0	0	0




## Scenario C: DB2 is working, it must be the network

### The setting:

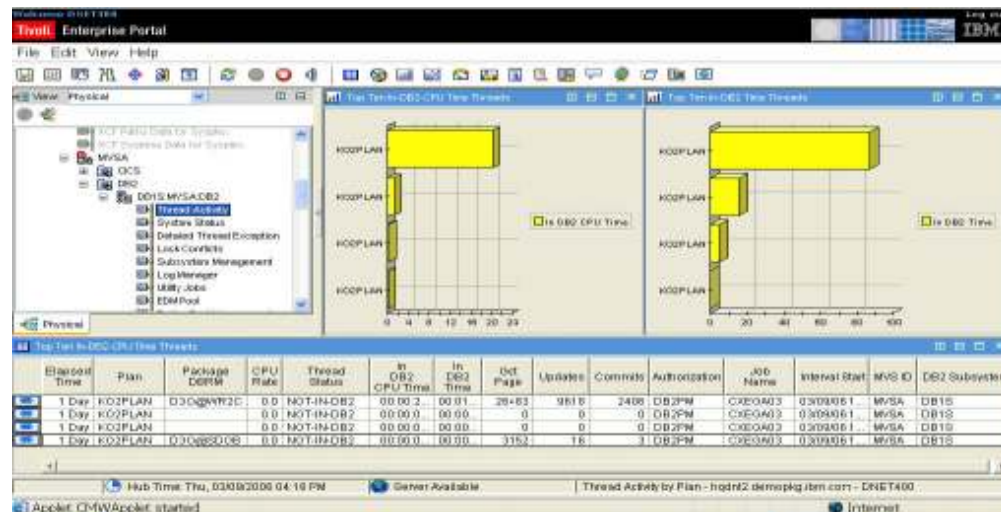
A multi-tier application framework is being used by a team of programmers to develop a Java application. The application is stored as large binary objects (BLOBs) in a DB2 on z/OS database. Each programmer retrieves, changes, and then saves a BLOB. Long delays that occur sporadically during the save are frustrating the application team.

## DB2 is working, it must be the network ...




## DB2 is working, it must be the network ...

1. Facing revolt from his team, the team leader asks the DB2 systems programmer to check for performance problems.
2. The DB2 systems programmer checks thread CPU time, lock contention, and query plan, among other things. He determines that DB2 is not the cause of the slowdown.



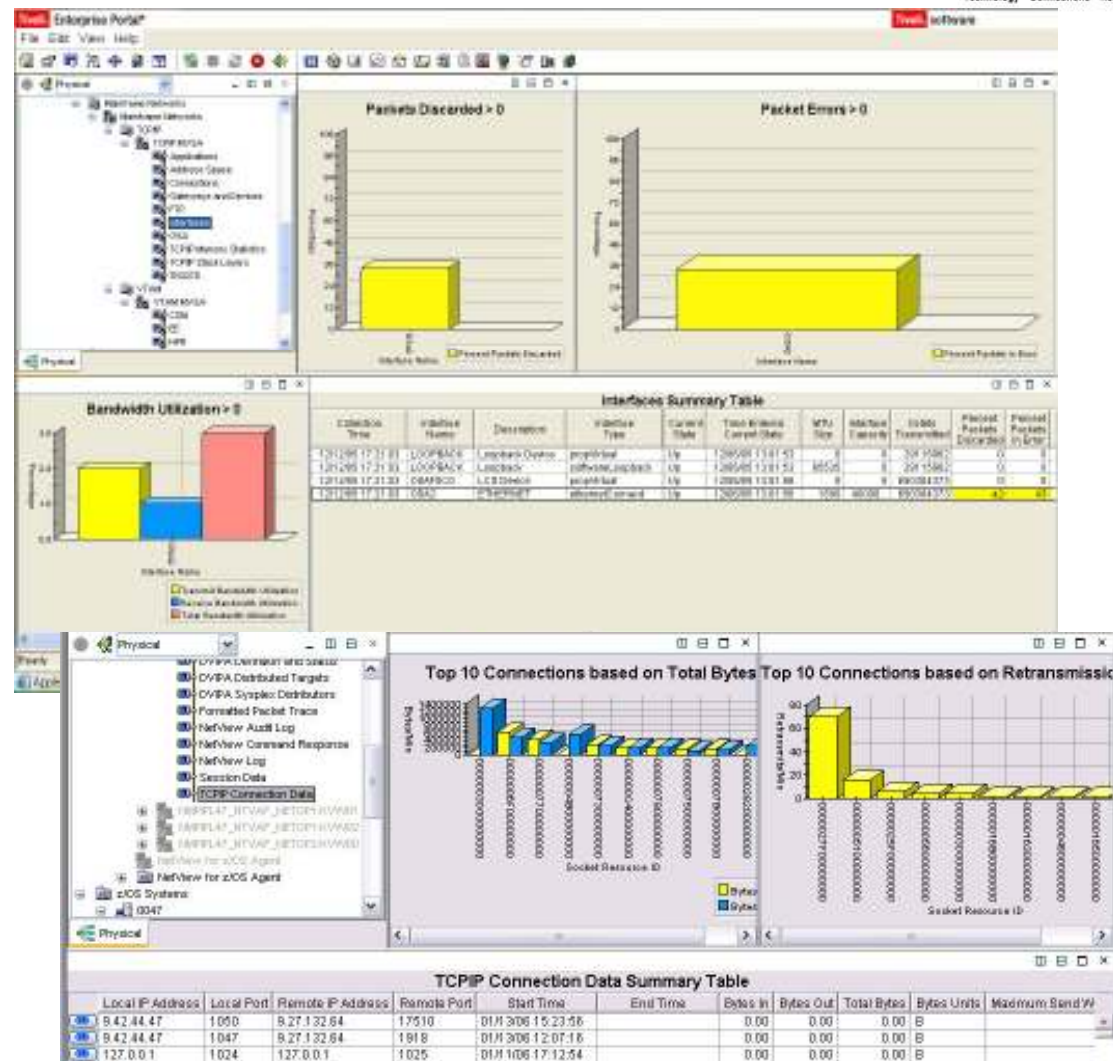




**SHARE**  
Technology • Connections • Results

## DB2 is working, it must be the network ...

3. Expecting that the problem may be due to an underlying network problem, the team leader turns to Johann for help.
4. Johann views the DB2 application and associated connections. Large amounts of data is being transferred over the DB2 connections with no retransmits or out of order segments.

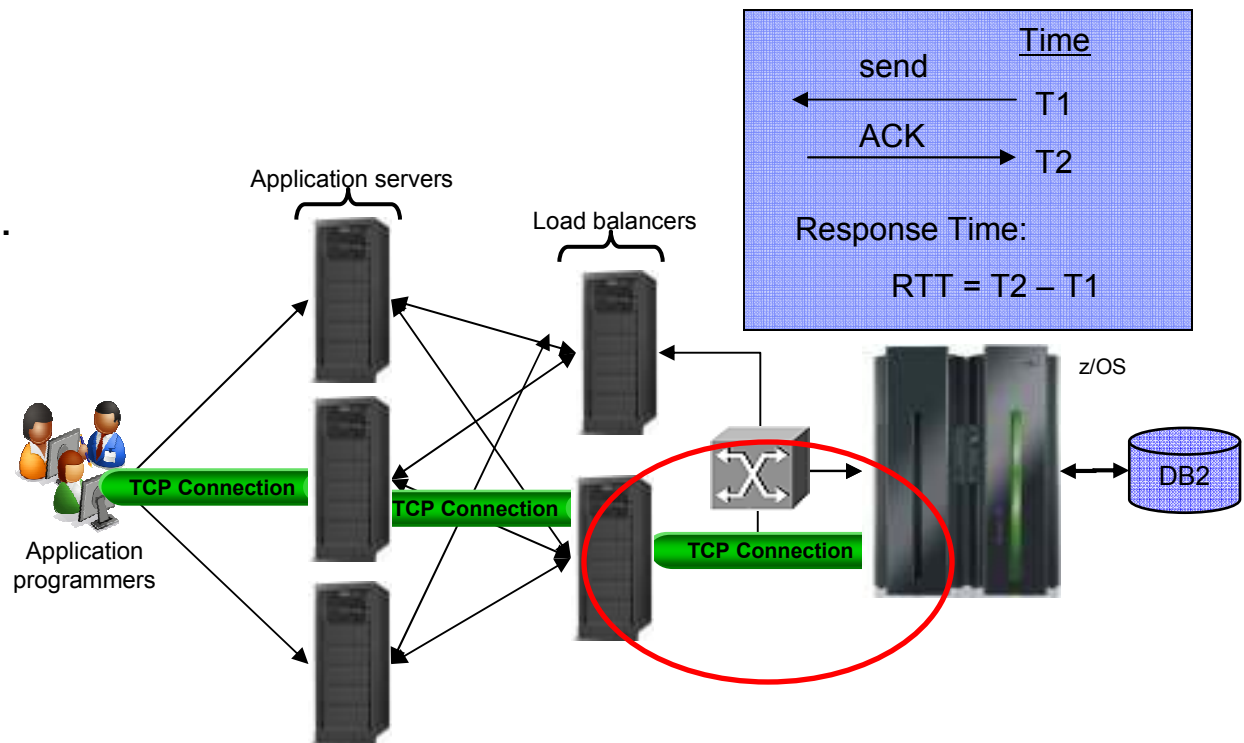




## DB2 is working, it must be the network ...

5. Interesting... Response time and response time variance are higher than expected (0.5+ sec, 0.5+). Also, much more data is being sent from DB2 than received from the remote system.

Why is ACK from remote system taking so long?

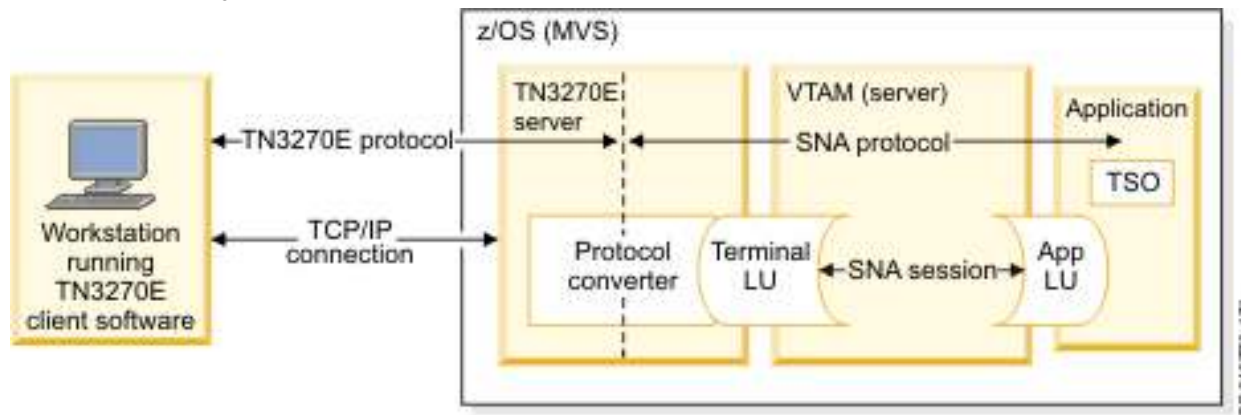


6. Working with distributed network and other SMEs to identify and resolve.

## Scenario D: Erratic response times for TN3270 application

The setting:


Users are becoming frustrated at response times with an SNA application. All access is through TN3270. The response times are on average very fast, but vary widely over the course of a day.






## Erratic TN3270 response times ...

1. A user opens a trouble ticket. Annette contacts the user who identifies a TN3270 session (TCP00072) that exhibits the erratic behavior.




2. The average response time and average SNA response time are fairly high. In contrast, average IP response time is good, so does not appear to be a network problem.



## Erratic TN3270 response times ...

3. Looking further, the bucket counts show that there have been a number of transactions with poor response time and a number with good response times but not much in between.
4. Annette passes the problem to the SNA application support team, which identifies and resolves the issue.



### High average SNA response time? Investigate:

- High application workload spike
- z/OS system resource constraints.




## Scenario E: Application Performance Problem

The setting:


A company is starting to protect more and more of its IP traffic using encryption. The deployment has gone well and the IT operations staff is trained and ready. A user calls the help desk because a file transfer is taking a long time.

# z/OS IP Security Support



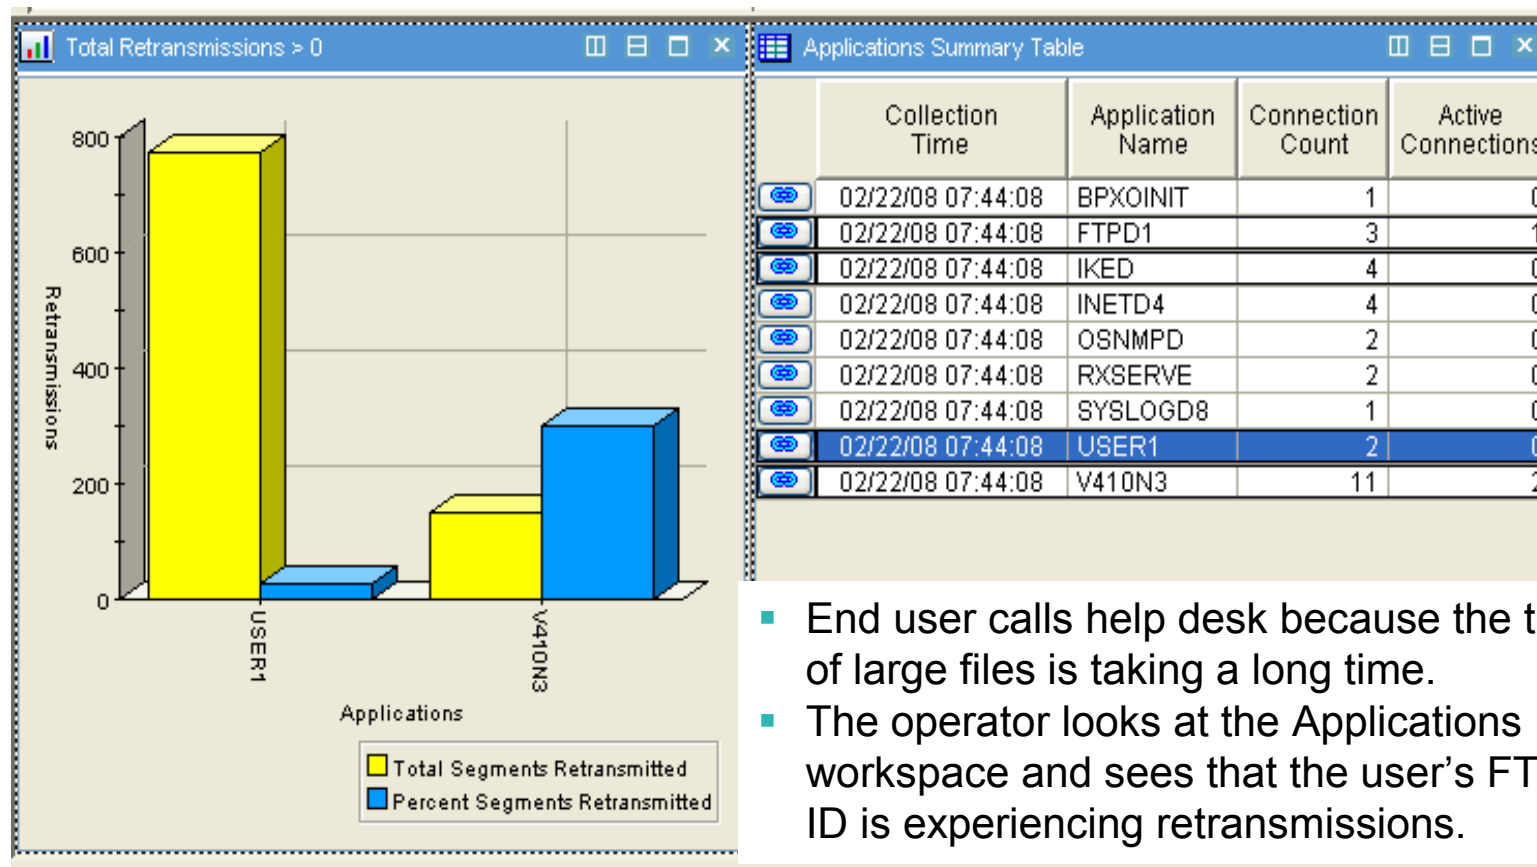
## Common Problems and Symptoms:


- Filter added in wrong order
  - Loss of connectivity to applications
- Security policies at endpoints are incompatible
  - Loss of connectivity to applications
  - Tunnel activation failures
- Loss of network connectivity between security endpoints
  - Loss of connectivity to applications
  - Tunnel activation failures
- Cryptographic services unavailable, misconfigured, or insufficient
  - Application performance is slow
  - Loss of connectivity to applications
  - Tunnel activation failures



**SHARE**  
Technology • Connections • Results

## Application performance problem ...

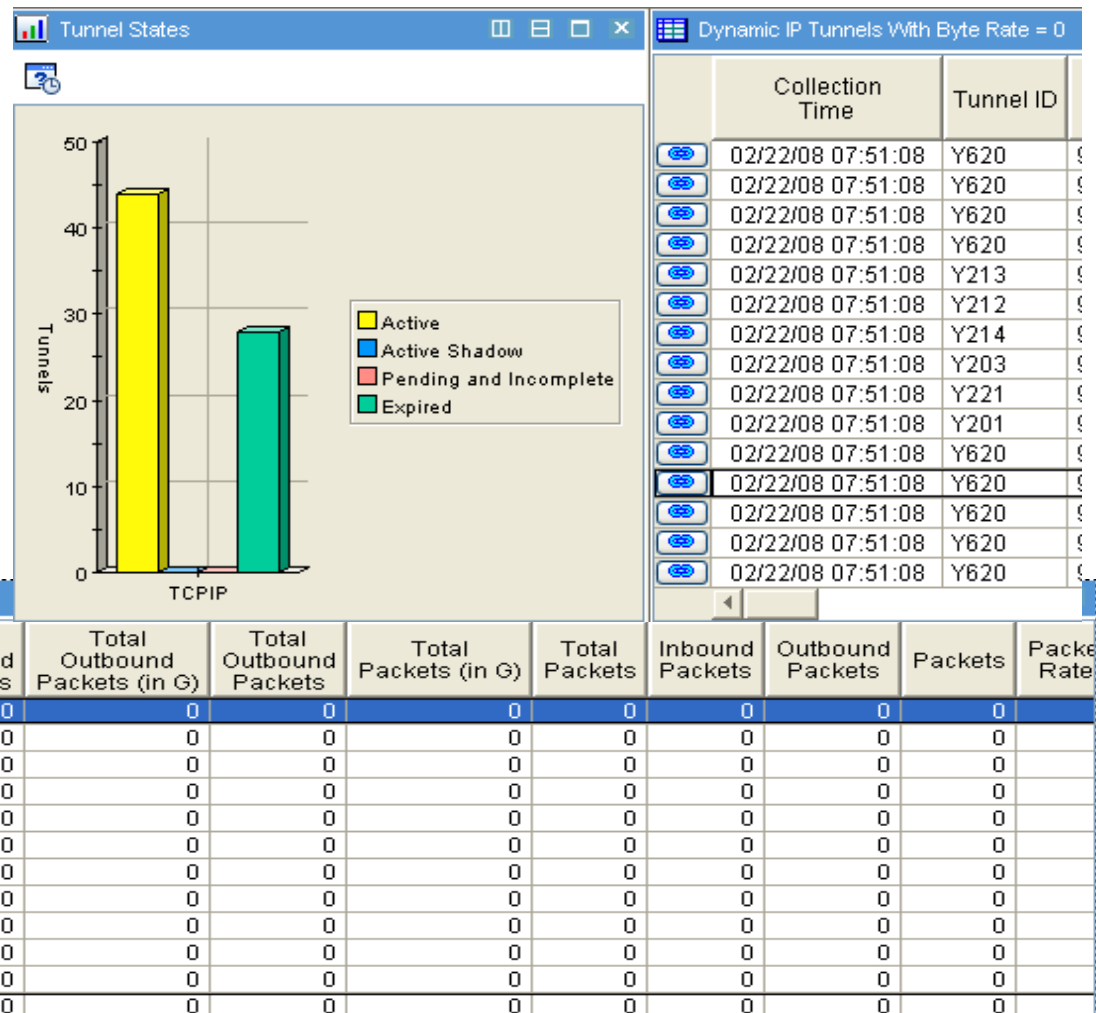




**SHARE**  
Technology • Connections • Results

## Application performance problem ...

- The systems programmer finds the IP filters for src/dst IP address, then finds the associated dynamic tunnels
- There are a high number of expired tunnels.
- The tunnel associated with the user's transfer has data rates of 0 and there are many tunnels with the same tunnel ID indicating it has been refreshed many times.







## Application performance problem ...

- The systems programmer examines the tunnel refresh and expiration information.
- The tunnel is being refreshed every 2 to 10 seconds.
- The systems programmer corrects the refresh time for the tunnel, which fixes the performance problem.

Dynamic IP Tunnels By Tunnel ID							
ID	Current Life Size	Life Size	Refresh Life Size	Life Expiration Time	Life Refresh Time	VPN Life Expiration Time	Activation Method
	0	32768	26892	02/22/08 07:52:08	02/22/08 07:51:58	02/23/08 07:44:13	ONDEMAND
	0	32768	24982	02/22/08 07:52:06	02/22/08 07:51:52	02/23/08 07:44:13	ONDEMAND
	0	32768	26249	02/22/08 07:52:05	02/22/08 07:51:54	02/23/08 07:44:13	ONDEMAND
	0	32768	26810	02/22/08 07:52:04	02/22/08 07:51:54	02/23/08 07:44:13	ONDEMAND
	0	32768	26125	02/22/08 07:52:02	02/22/08 07:51:50	02/23/08 07:44:13	ONDEMAND
	0	32768	27834	02/22/08 07:52:01	02/22/08 07:51:52	02/23/08 07:44:13	ONDEMAND
	0	32768	23507	02/22/08 07:51:59	02/22/08 07:51:43	02/23/08 07:44:13	ONDEMAND
	0	32768	24314	02/22/08 07:51:58	02/22/08 07:51:43	02/23/08 07:44:13	ONDEMAND
	0	32768	27770	02/22/08 07:51:56	02/22/08 07:51:47	02/23/08 07:44:13	ONDEMAND
	0	32768	25679	02/22/08 07:51:54	02/22/08 07:51:42	02/23/08 07:44:13	ONDEMAND
	0	32768	26865	02/22/08 07:51:52	02/22/08 07:51:42	02/23/08 07:44:13	ONDEMAND
	0	32768	27671	02/22/08 07:51:51	02/22/08 07:51:42	02/23/08 07:44:13	ONDEMAND





## Questions?

# Identifying and Solving Network Performance Problems on zEnterprise

Mac Holloway (mhollowa@us.ibm.com)  
Dean Butler (butlerde@us.ibm.com)



Tivoli Software

# Backup



