

Identifying and Solving Network Performance Problems on zEnterprise

Dean Butler (butlerde@us.ibm.com)



Context Chart



IBM's Integrated Service Management (ISM) framework can optimize costs and streamline operations

This session is focused on:

Integrated Service Management



VISIBILITY



See your business services

Understand health and performance of services across your enterprise infrastructure

CONTROL



Manage service risk and compliance

Govern and secure complex infrastructure and ensure regulatory compliance

AUTOMATION



Optimize business service delivery

Drive down cost, minimize human error and increase productivity

Agenda

- Is this familiar?

“A critical application is ‘broken’. We all get on a bridge line. My peers say ‘My 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.”
- Knowing the problem space
 - Configuration issues
 - Operational issues
- zEnterprise ... Integrated Service Management
- Scenarios

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

- SNMP
 - Is it configured (SNMPD.CONF)?
 - Is OSNMP job running?
 - Are you using the right IP address (loopback address, 127.0.0.1)?
 - Do you have the right community name (check SNMPD.CONF)?
 - Are you using the right port (default is 161)?
 - SNMP requests are timing out?

- IPsec
 - IKE daemon started?
 - PAGENT daemon started?
 - IPsec Network Management Interface (NMI) access authorized?

- TN3270 & FTP
 - z/OS Communications Server real-time SMF NMI enabled?
 - Monitoring app authorized?
 - Sliding window or bucket count data – configured in Telnet server profile?

Configuration Issues – examples ...

- SNA NMI
 - z/OS Communications Server SNA NMI enabled?
 - OMVS segment created for VTAM?
 - Monitoring app authorized?

- VTAM
 - Is major node running?
 - Is SNA data collection configured?
 - Is your monitoring app in the VTAMLST?
 - Is the PMI exit available to VTAM?

Configuration Issues – an approach

The screenshot shows the 'Agent Status' window with several tables and navigation elements. Red circles highlight the 'Physical' view button, the 'Agent Status' table, the 'SNA Collector Status' table, and the 'TCP Collector Status' table.

Time	TCP Collection Interval	TCP Collector SNMP Parameter Dataset Name	PAGENT Daemon Started	IKE Daemon Started	SNA Collection Started	SNA Collection Start Time	SNA Collection Interval	Virtual IO Unit Name
19:53	1	USER.PARMLIB(KN3SNMP)	Yes	Yes	Yes	04/22/10 13:19:54	1	VIO

Collection Time	Agent VTAM Major Node Name	Agent VTAM Major Node Status	Agent VTAM Application Name	Agent VTAM Application Status	PMI Exit Name	PMI Exit Status	SNA NMI Enabled	EE And HPR Collection	ALL HPR Collection	CSM Buffer Reporting Collection	Buffer Pool And VTAM Environment Collection
04/22/10 15:34:40	CTDN3N	ACTIV	CTDN3SP	ACTIV	KN3AMV00	ACTIVE	Yes	Yes	No	Yes	Yes

Time	TCP/IP Profile Dataset Name	Host Name	IP Address	SNMP Agent Jobname	SNMP Agent Port	SNMP Version	SMF Service Enabled	IPv4 Security Enabled	IPv6 Security Enabled	Connections And Applications Collection	IP Security Collection	Routing Table Collection	Routing Table Collection Frequency
	USER.PARMLIB(TCPPROF2)	TVT4085	9.42.32.85	OSNMPD	161	snmp...	Yes	Yes	Yes	Yes	Yes	Yes	

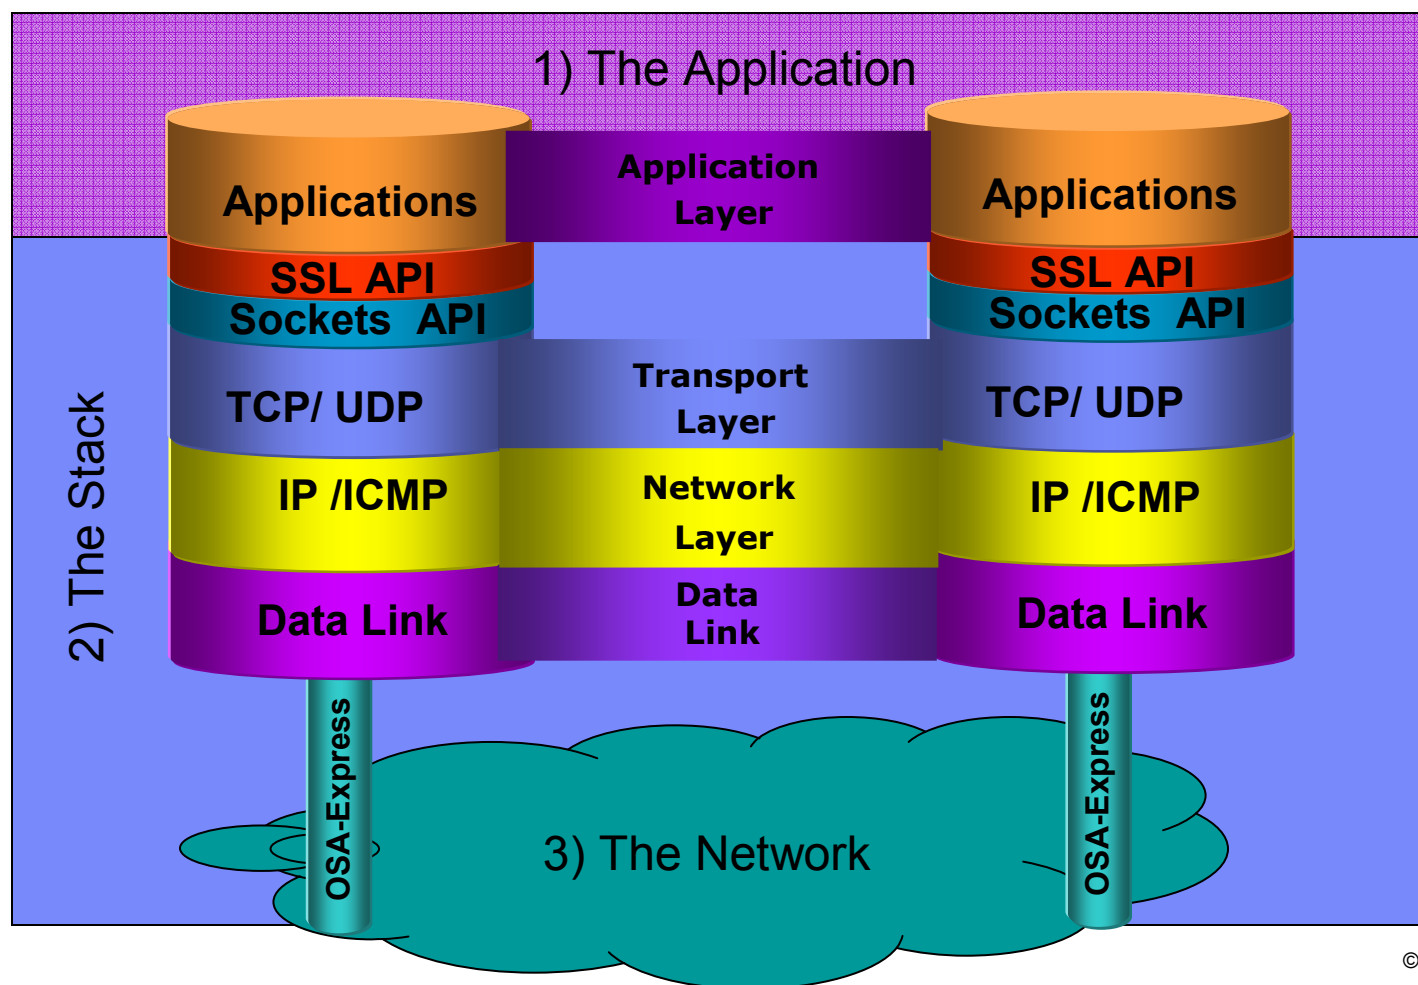
Status on:

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

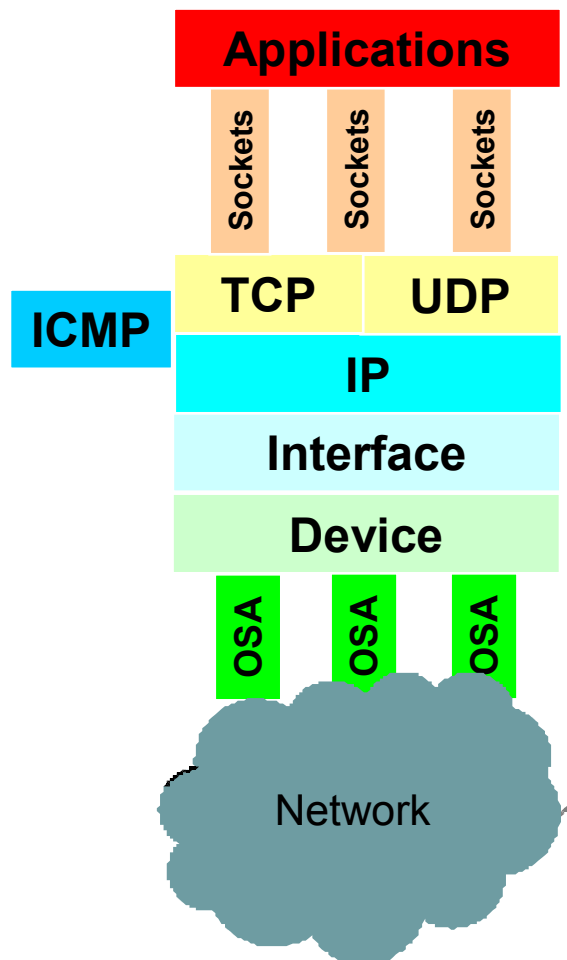
Situations/actions based on status

Operational Issues

- Divide into 3 “locations”
- Problem sources



TCP/IP Key Metrics



- **TCP/IP Error Handling**

- UDP, IP, and interfaces discard transmission units
- TCP retransmits unacknowledged segments and drops the connection if retransmission threshold is reached
- Applications that use UDP may retransmit data

- **Interfaces key metrics**

- Status– Is the interface up or down?
- Discards
 - Inbound– indicates possible problem w/ interface, adjacent node, or network
 - Outbound– indicates possible problem w/ interface, IP stack, or application
- Throughput-- uneven distribution across defined interfaces indicates possible problems with path to gateway or lack of network availability


- **IP key metrics**

- Discards
 - Inbound– checksum errors can be caused by problems w/ local interface or problem at adjacent node; security definition problems; routing errors
 - Outbound– security definition problems; routing errors
- Throughput
 - Significant drop for extended time can indicate network problem
 - Significant increase for extended time can indicate application error.
- Fragmentation/Reassembly: frequent fragmentation/reassembly can indicate an application problem or an error in a recent network configuration change

- **UDP key metrics**

- Discards– checksum errors or no port
- Throughput

TCP/IP Key Metrics (cont)

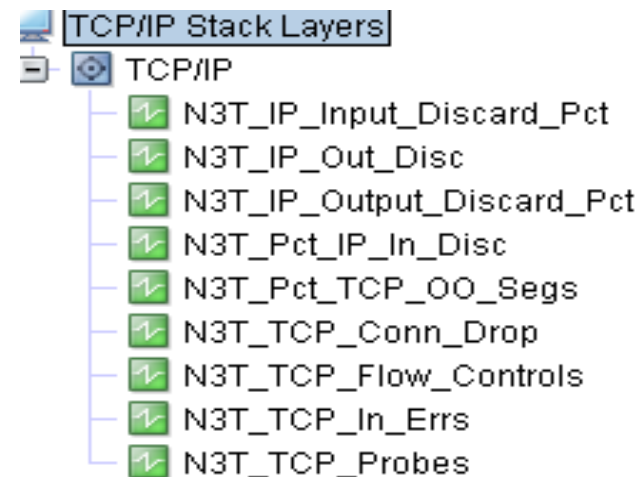
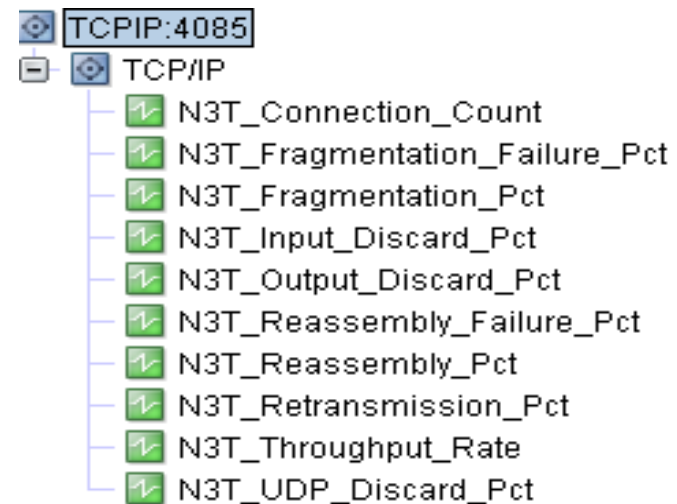
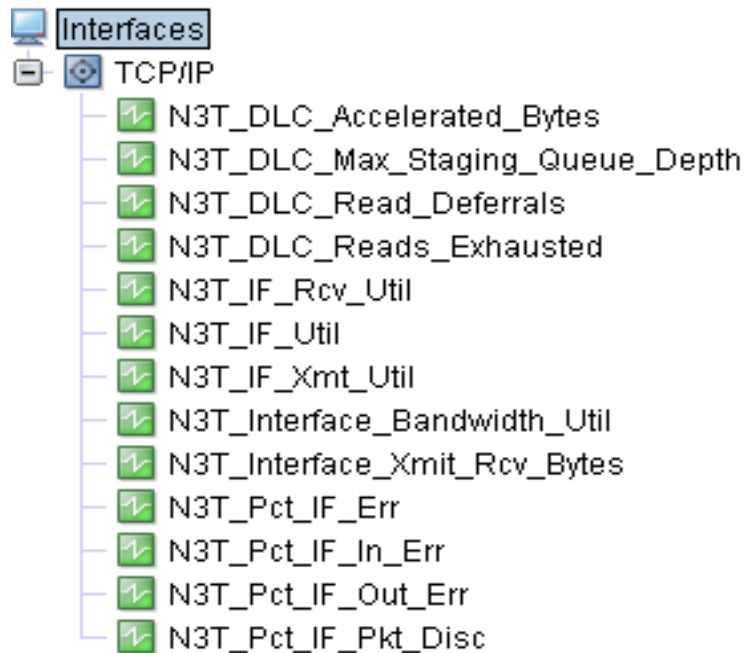


▪ TCP key metrics

- Session count– large change from typical value for a significant amount of time can indicate a problem w/ application, OS, or network
- Connections dropped– large number during one or more consecutive intervals can indicate problems w/ application, OS, or network
- Retransmits-- ACKs for data sent not being received
 - Is data being sent successfully and arriving at partner node?
 - Is partner generating ACKs and successfully sending them?
- Duplicate ACKs– Data being sent is not being received by partner
 - Need to figure out whether data is being sent successfully, if so, then where is it being dropped?
- Segment errors received – checksum errors indicate problem w/partner endpoint
- Window probes– large numbers can indicate problem w/remote application
 - Partner has closed window (sent a 0 window) and no data can be sent
 - Window probe requests window be opened so data can be sent
 - If window probe threshold is reached connection is dropped

Monitor key metrics automatically using situations

- Metrics aggregated at the stack level are best high-level indicators for situations
- Historical data can be analyzed to determine the appropriate situation thresholds for your enterprise.

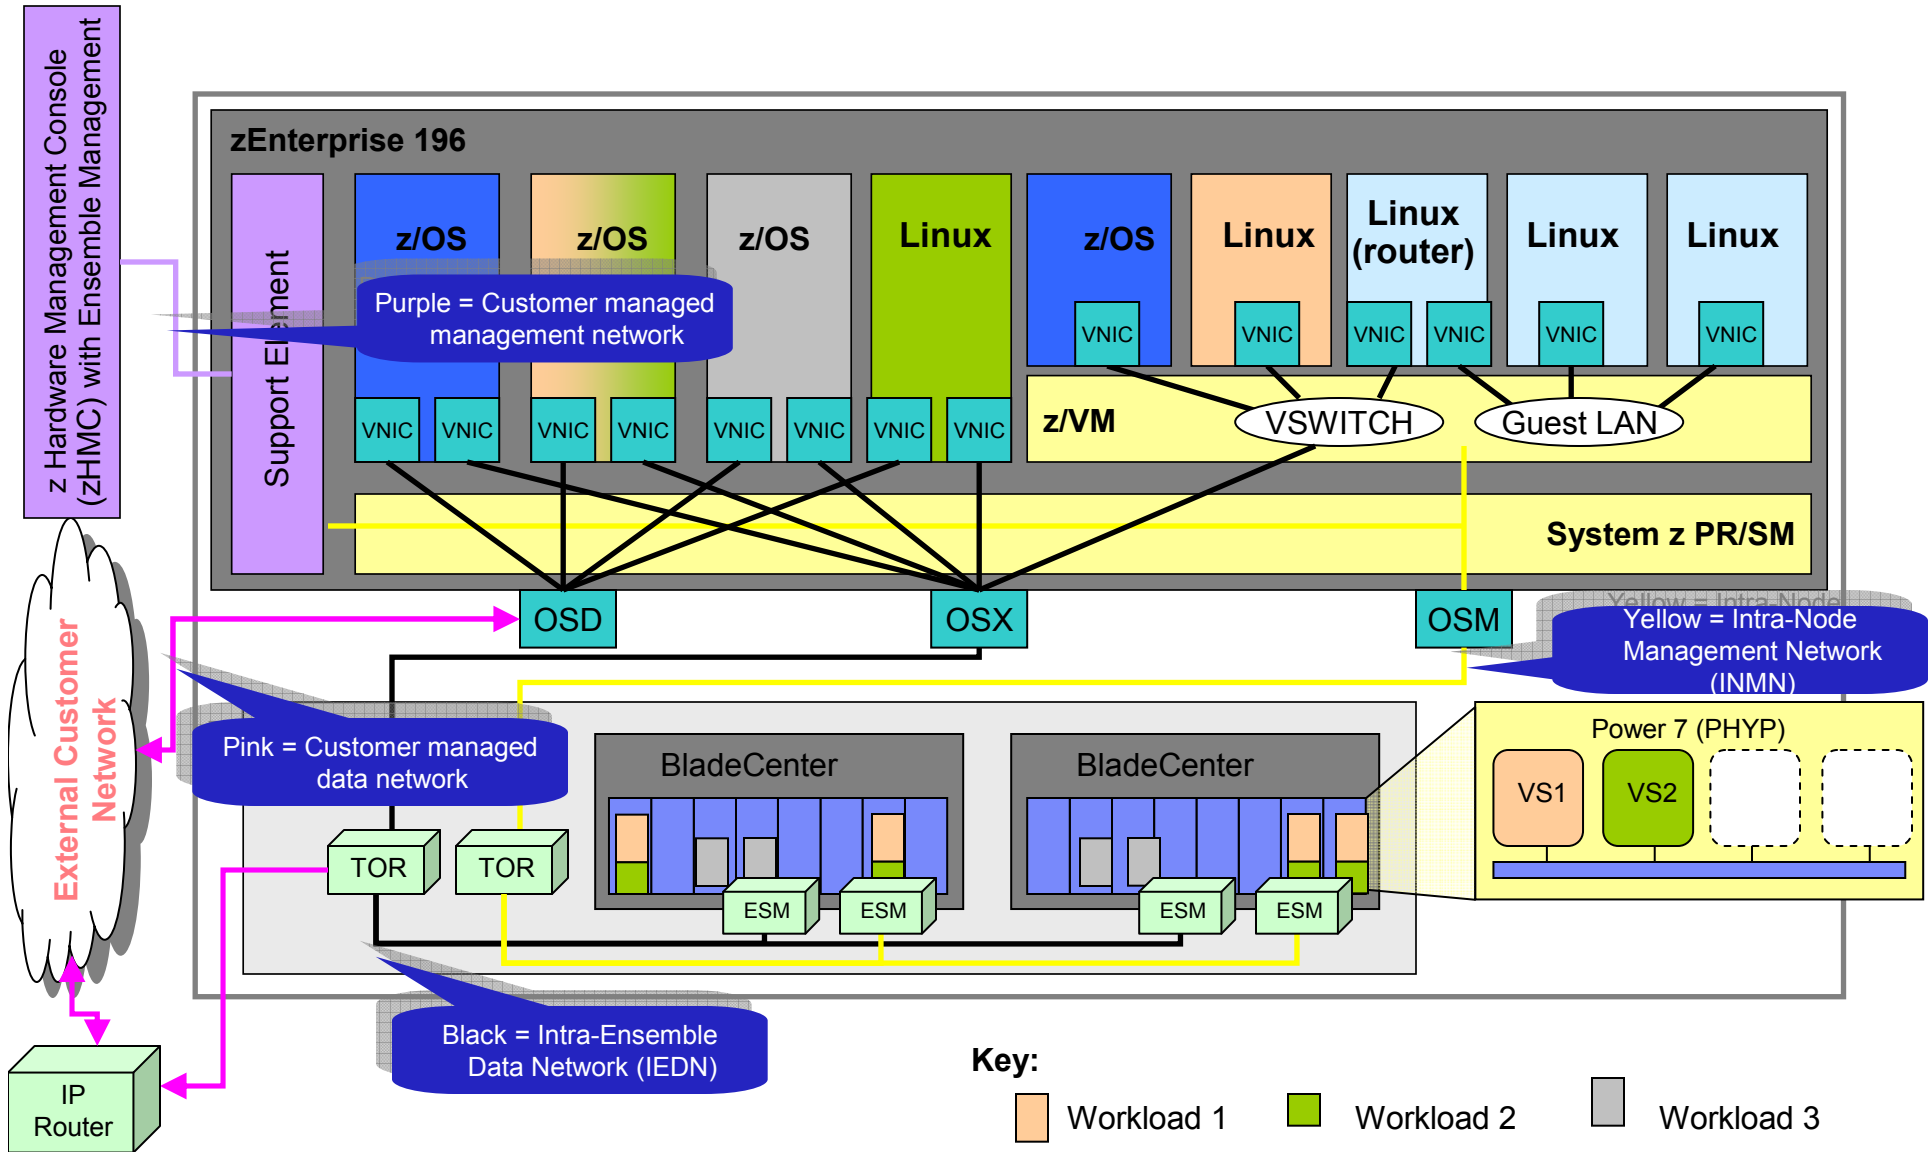


zEnterprise

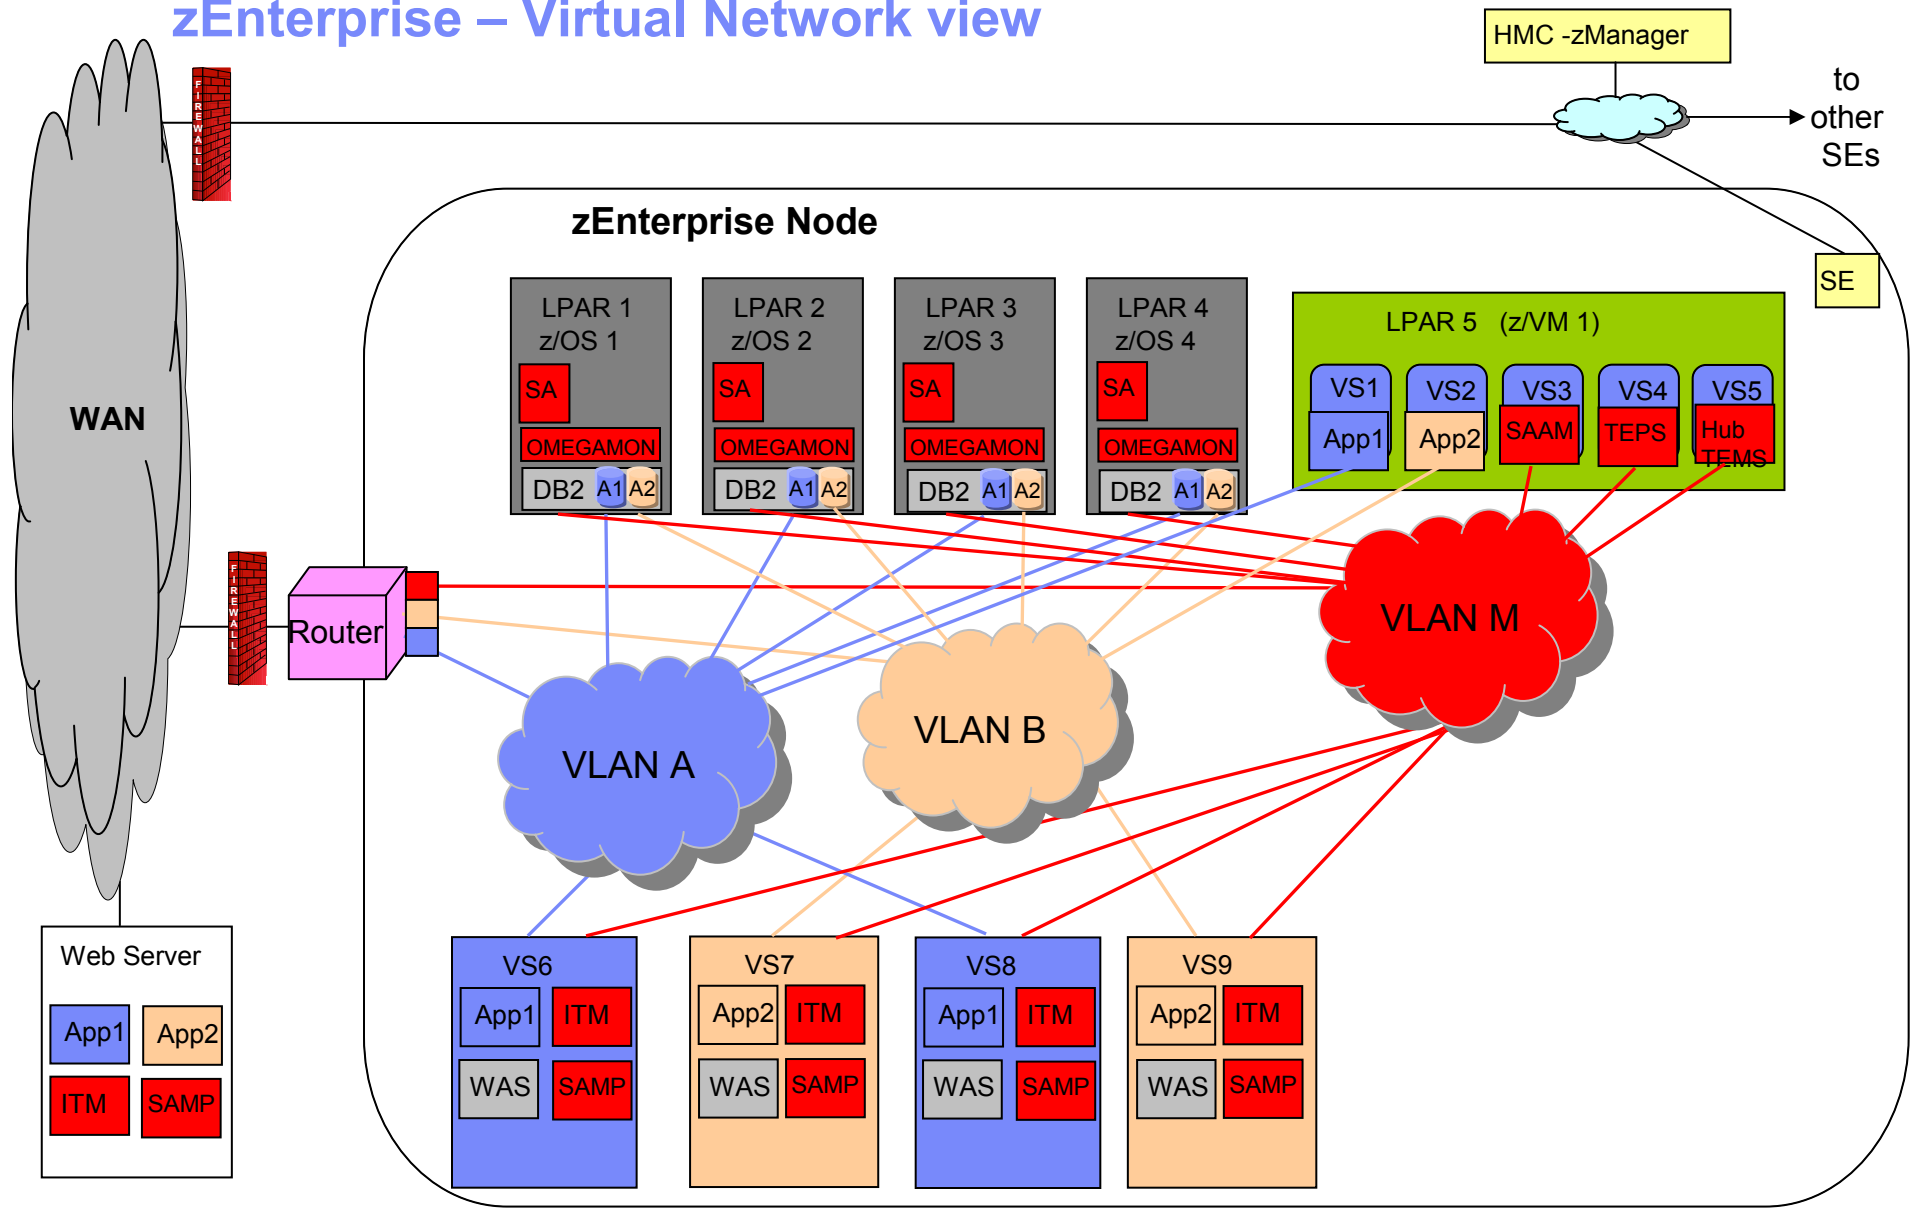
...

Integrated Service Management

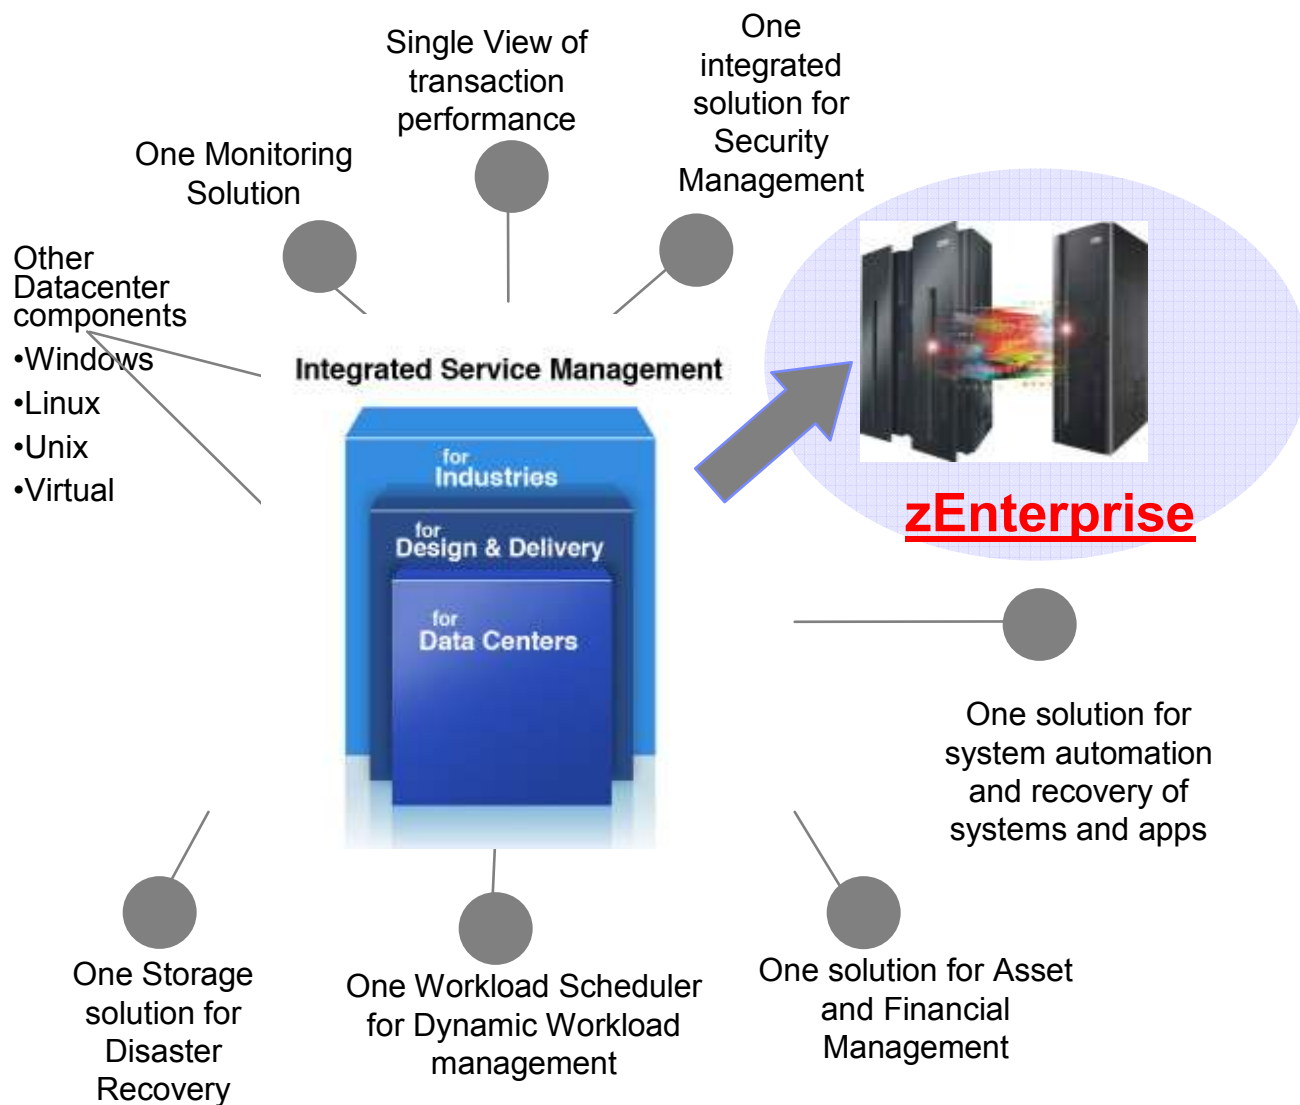
zEnterprise with zBX (z Blade Extension)



zEnterprise – Virtual Network view



ONE Service Management Engine - Leveraging the Strengths of ONE System



IBM Advantage

- ✓ *Visibility. Control. Automation™ with a SINGLE POINT OF CONTROL on a common infrastructure*
- ✓ *Consolidated view of the IT infrastructure*
- ✓ *Lower training and maintenance costs*
- ✓ *Integrated, consolidated reporting*
- ✓ *Asset optimization*


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 files were not updated. 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

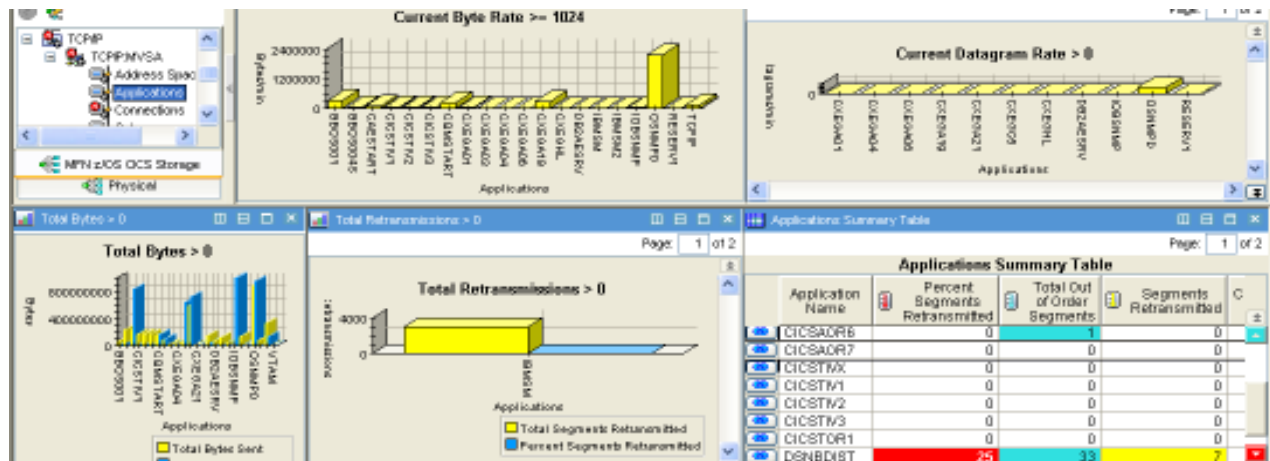
Collection Time	Application Name	FTP Type	Remote IP Address	Remote Port	Local IP Address	Local Port	User ID on Server	Client User ID	Session Start	Session End	Session Duration
08/11/08 23:52:22	FTPD1	Server	9.65.126.164	3000	9.42.45.179	21	USER2		08/11/08 23:52:18		0

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

Check Applications and Connections

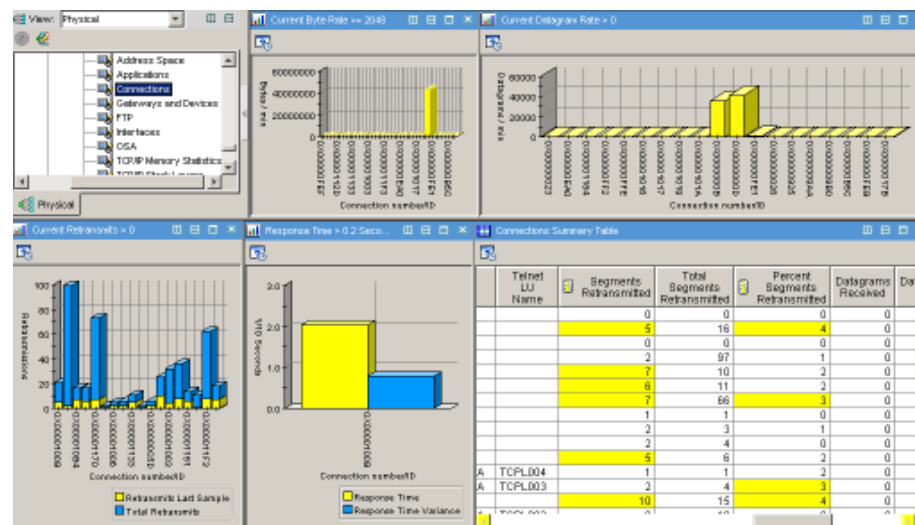
Applications:

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



Connections:

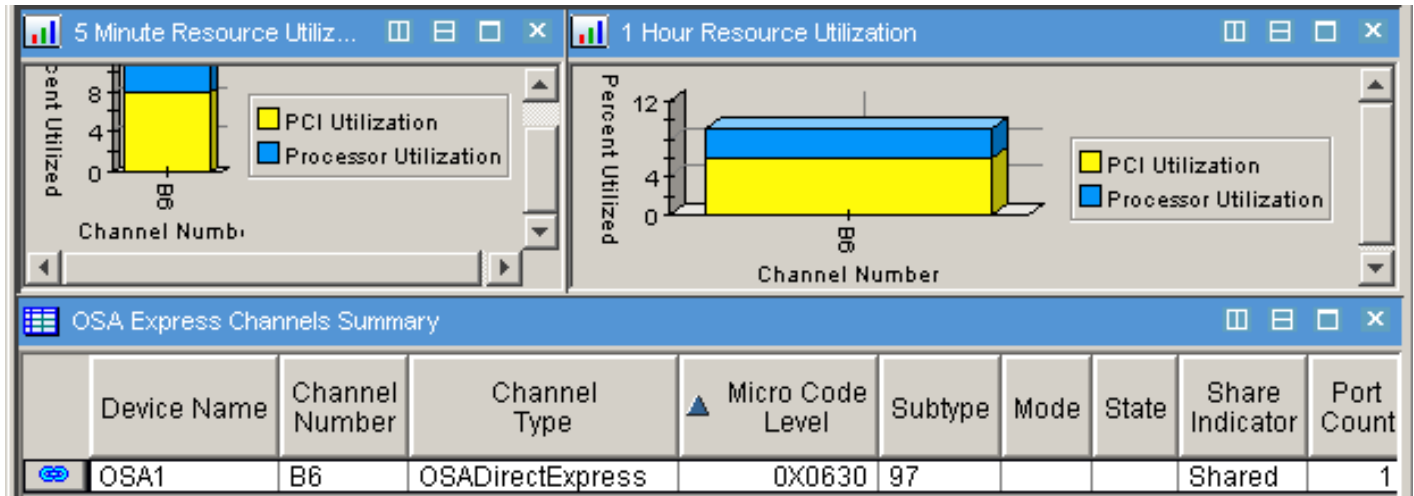
- Start time/duration
- Response Time
- Response Time Variance
- Retransmissions
- Transmit / Receive Rates
- Out of order segments



Check OSA and Interfaces

OSA

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



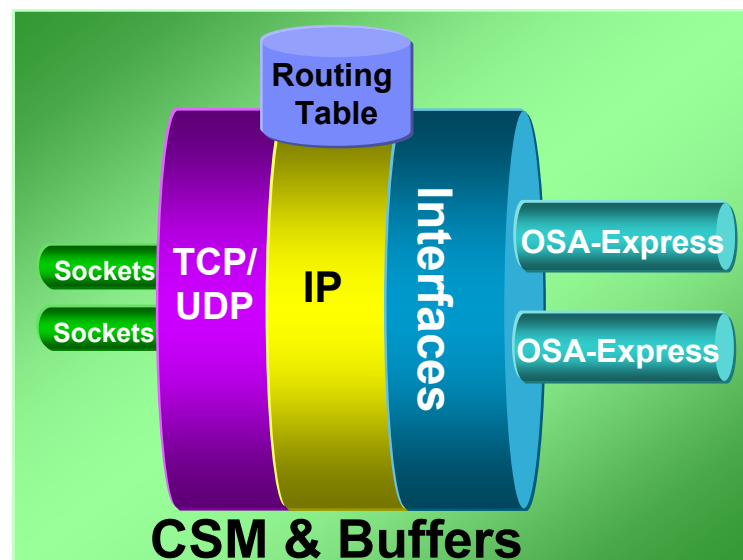
Interfaces

- 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
FAXCFSA	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?

The screenshot displays the 'FTP Sessions - IBM-4D3EDB432B0 - SYSADMIN' interface. It includes a Navigator pane on the left, two bar charts at the top right, and a summary table at the bottom.

Server Session Login Failures (Bar Chart):

User ID on Server	Login Failures
USER1	1.0
USER3	3.0
XXX	1.0
user1	1.0

FTP Client Session Count (Bar Chart):

FTP Client Sessions
1.0

FTP Server Session (Bar Chart):

FTP Server Session
1.0

FTP Session Summary Table:

User ID on Server	Client User ID	Session Start	Session End	Session Duration	Session End Reason Code	Session End Reason Description	Login Failure Reason Code	Login Failure Reason Description
USER3		08/12/08 00:11:56	08/12/08 00:11:56	0	0		1	Password_not_valid
USER3		08/12/08 00:12:16	08/12/08 00:12:16	0	0		1	Password_not_valid

Hub Time: Tue, 08/12/2008 01:03 AM | Server Available | FTP Sessions - IBM-4D3EDB432B0 - SYSADMIN

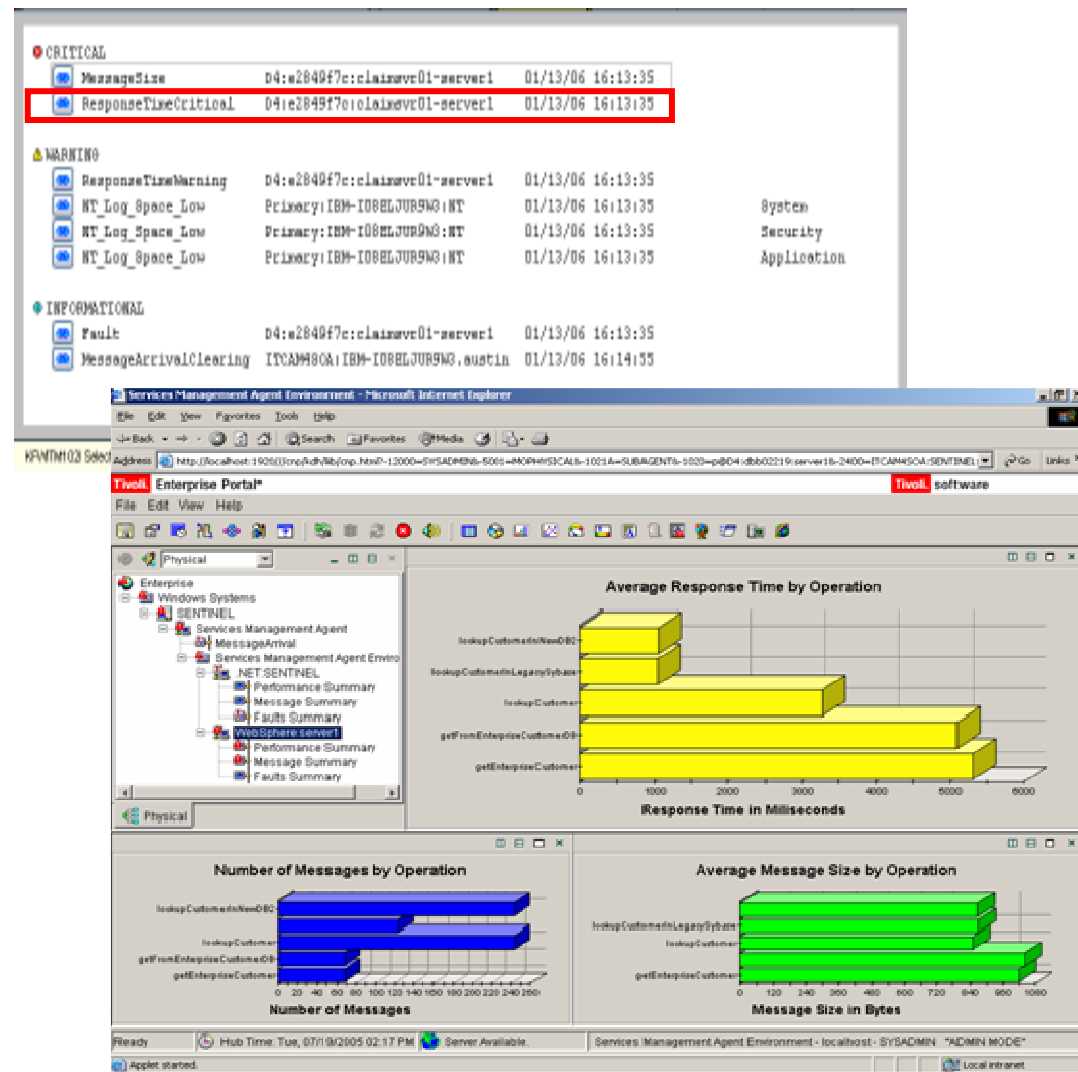
Scenario B: Slow response time in web service

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...

Slow response time in web service ...

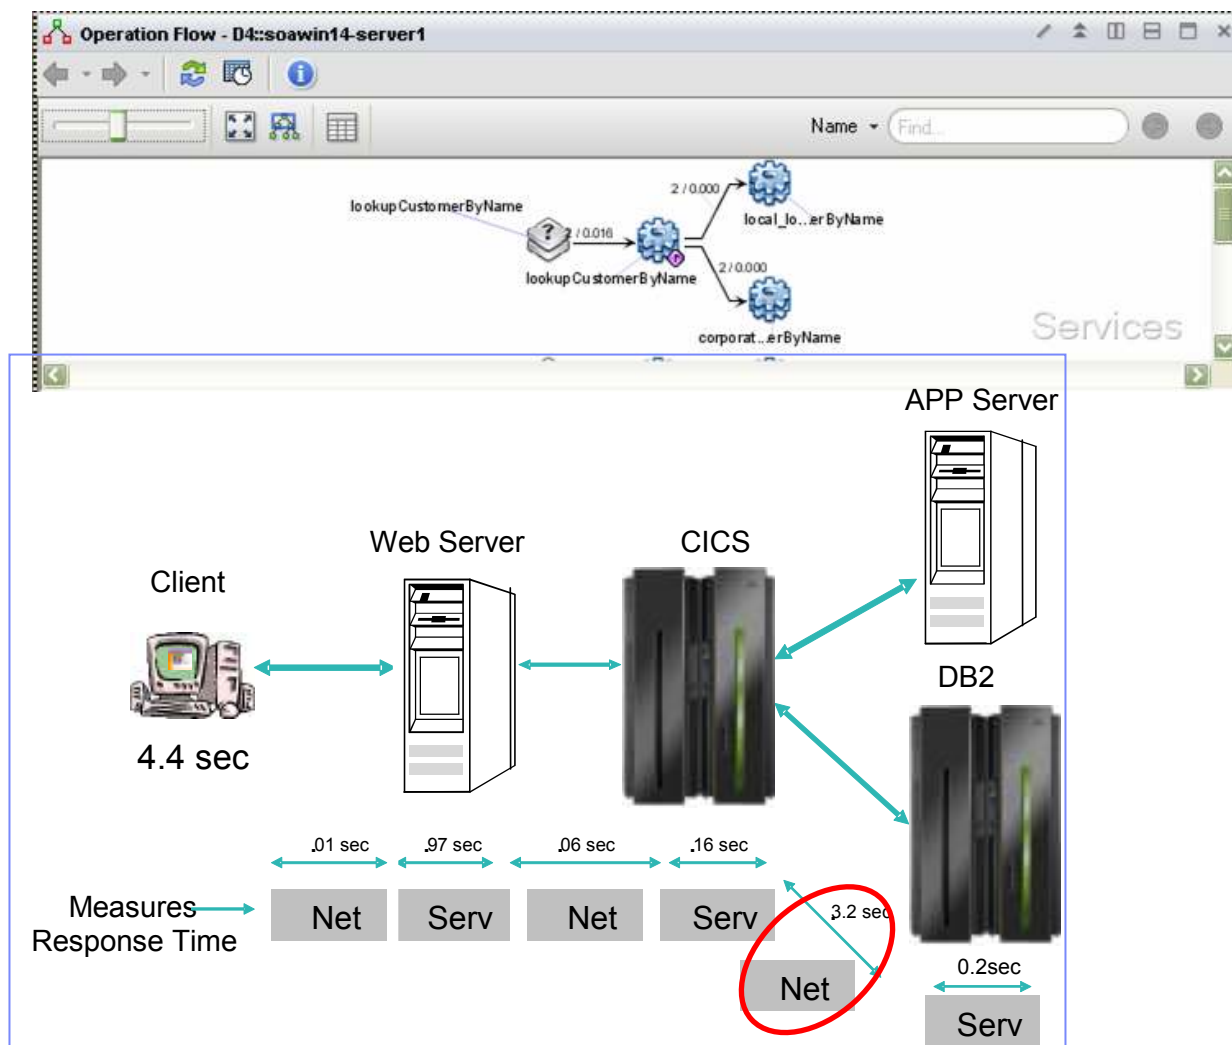
1. An alert identifies a response time problem. Annette, an operator, determines that slow response times are being recorded for the new web services.
2. Annette checks the number of requests and the message size and determines this is a normal volume of traffic. Annette passes the issue to Johann, a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

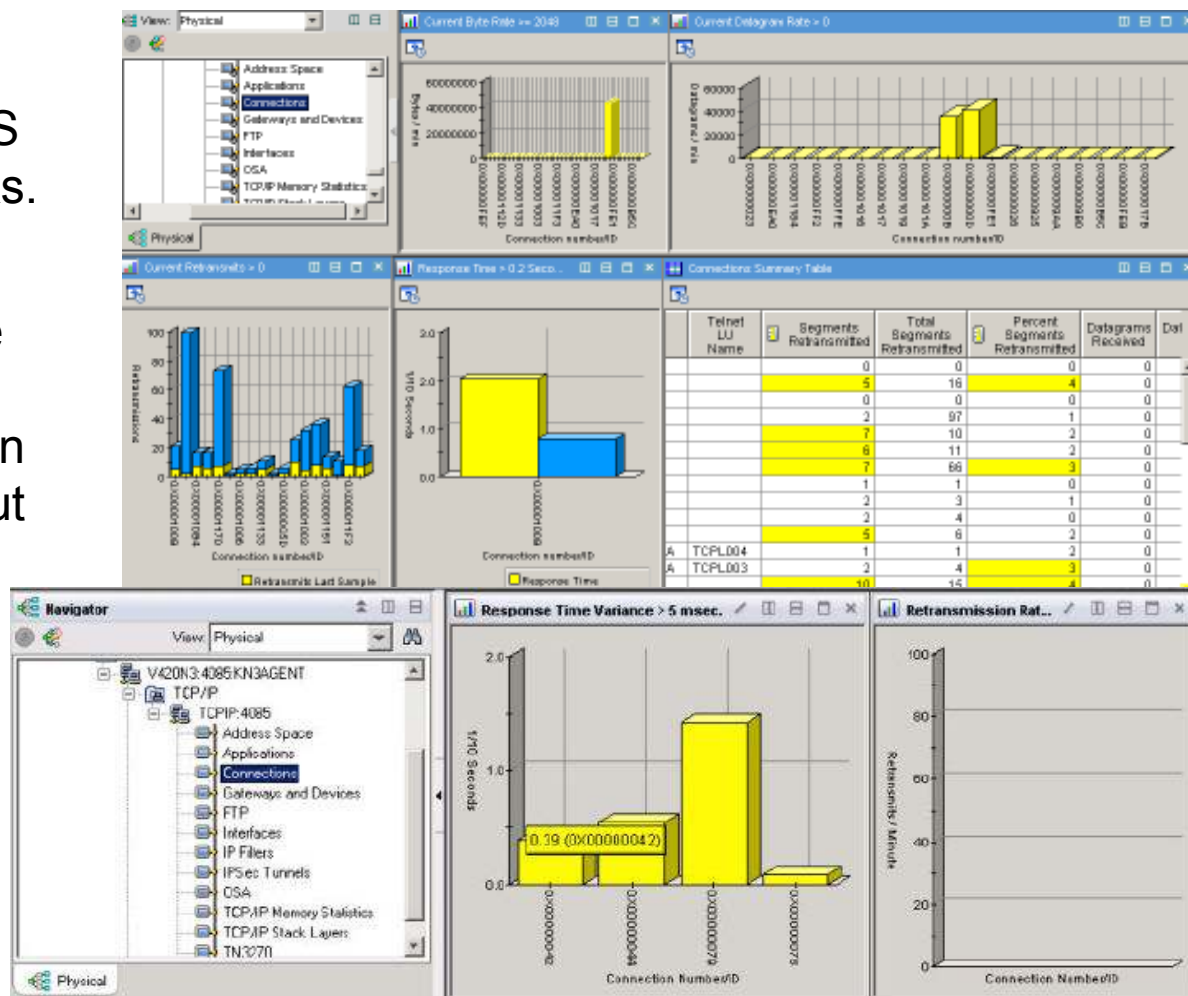
4. Problem appears to be with the network between the CICS and DB2 servers. These two LPARs are connected by a data center network.



Slow response time in web service ...

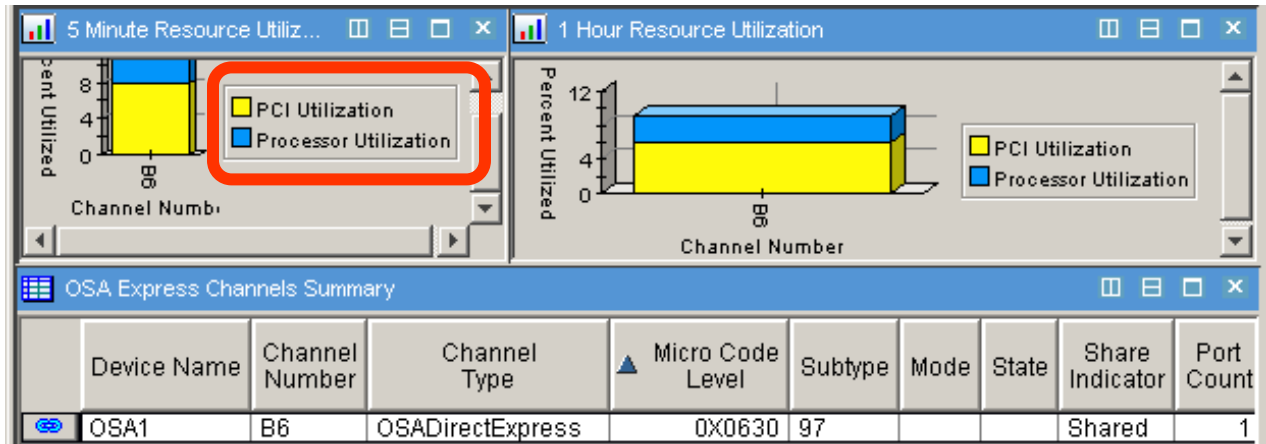
5. Johann views metrics for connections between CICS and DB2 on the two LPARs.

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.

The figure shows the 'OSA Express Ports Summary' table:

Collection Time	Port Name	Link Name	Port Type	Hardware State	Disabled Status	Service Mode	Configuration Name	Configuration Speed
12/16/08 13:34:34	OSA1	TCPIPLIN	oneThousandBaseTEthernet	enabled	0X0000	NotInServiceMode	IBM Default ConfigFile	1000Base T oneThousandMbFullDuplex

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

The figure shows the 'LPAR Usage Summary for Channel B6' table:


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 ...

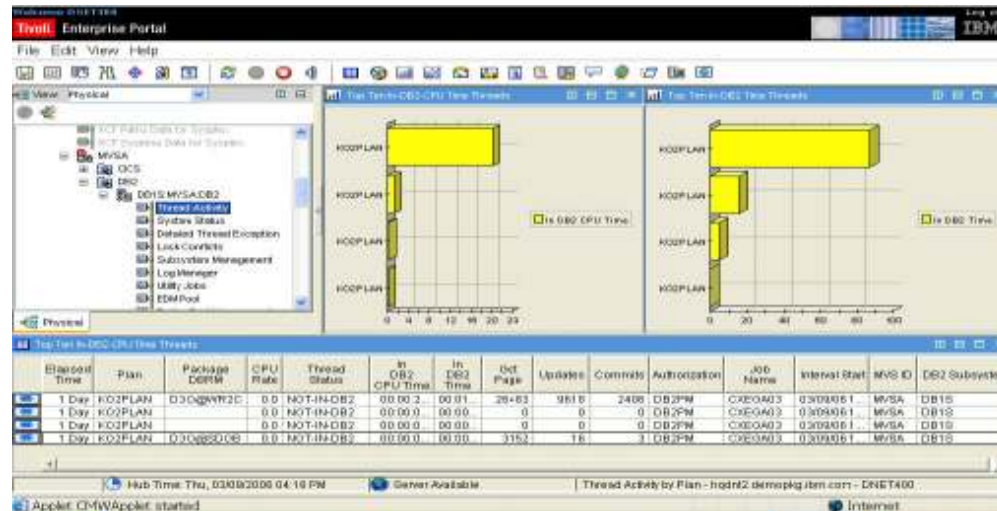


DB2 is working, it must be the network ...

1. Facing revolt from his team, the team leader asks the DB2 systems programmer to check for performance problems.



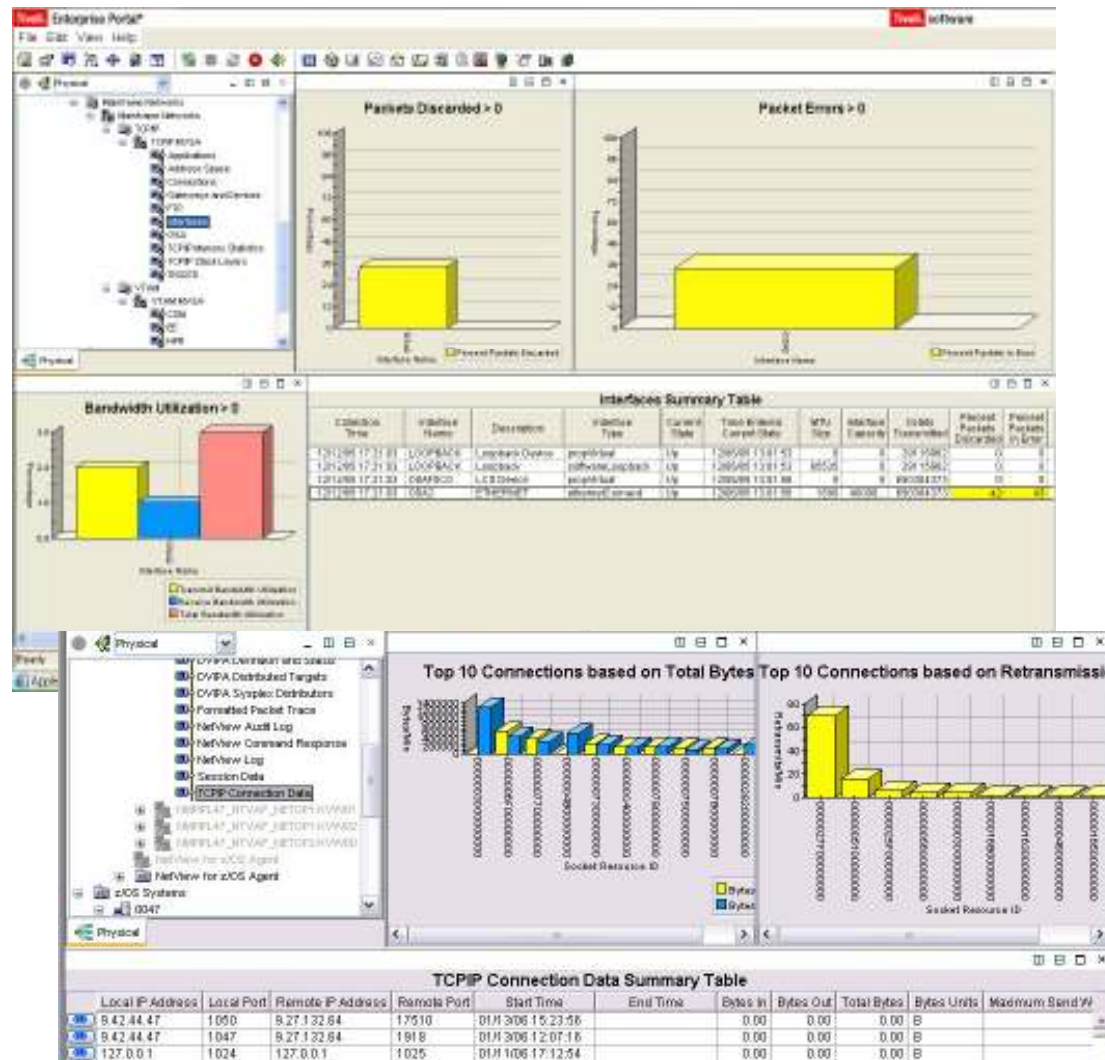
2. The DB2 systems programmer checks thread CPU time, lock contention, and query plan, among other things. He determines that DB2 is not the cause of the slowdown.



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.

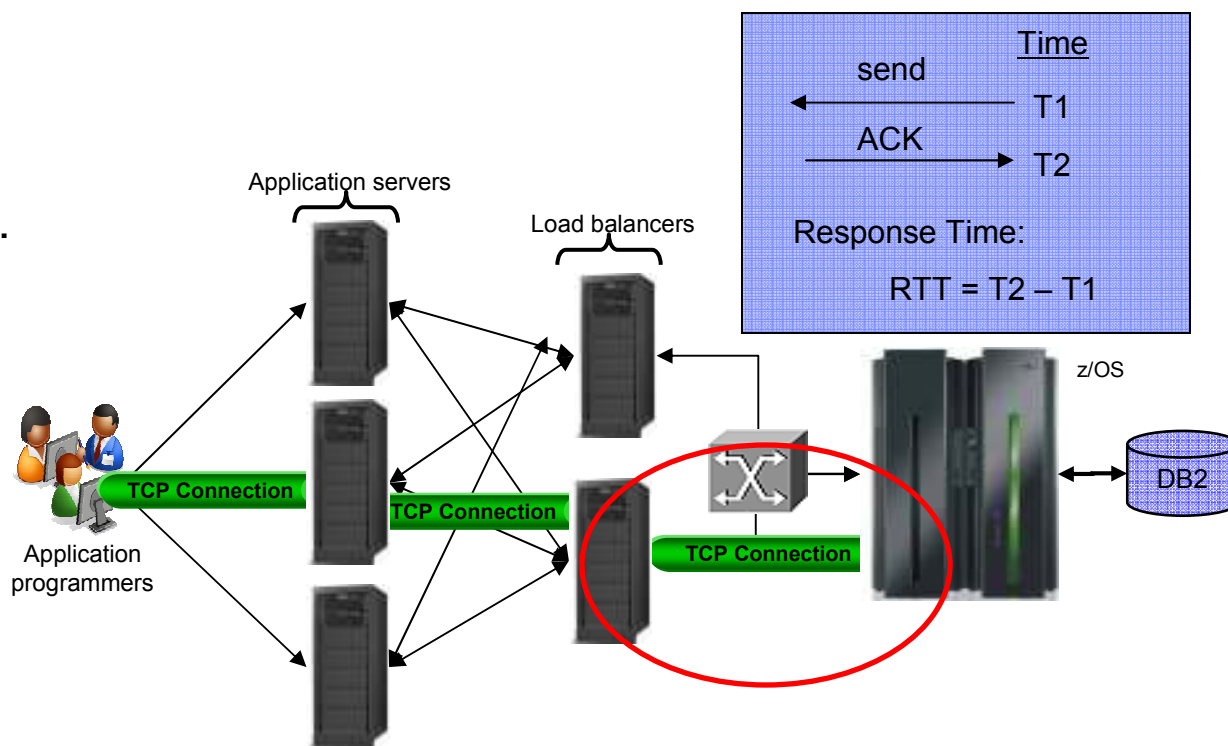
4. Johann views the DB2 application and associated connections. Large amounts of data is being transferred over the DB2 connections with no retransmits or out of order segments.



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?

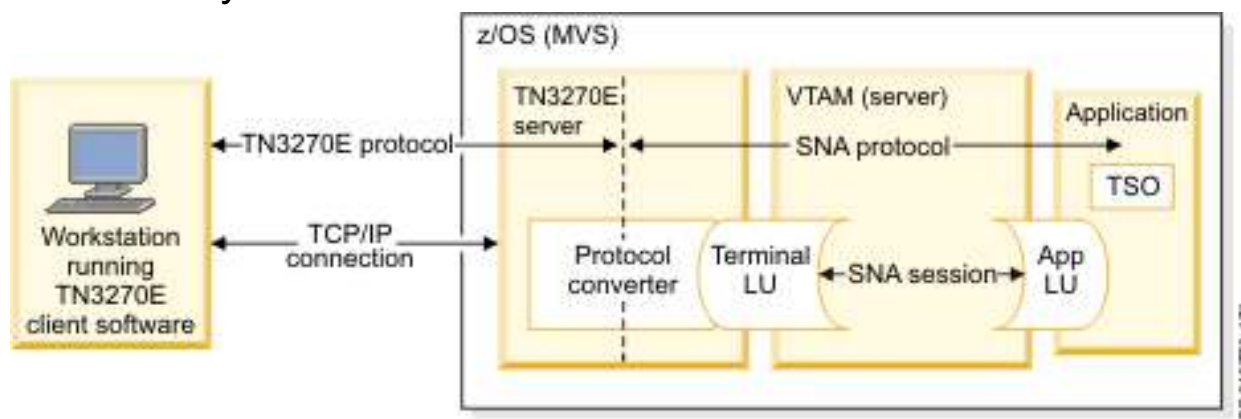


6. Working with distributed network and other SMEs to identify and resolve.

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.

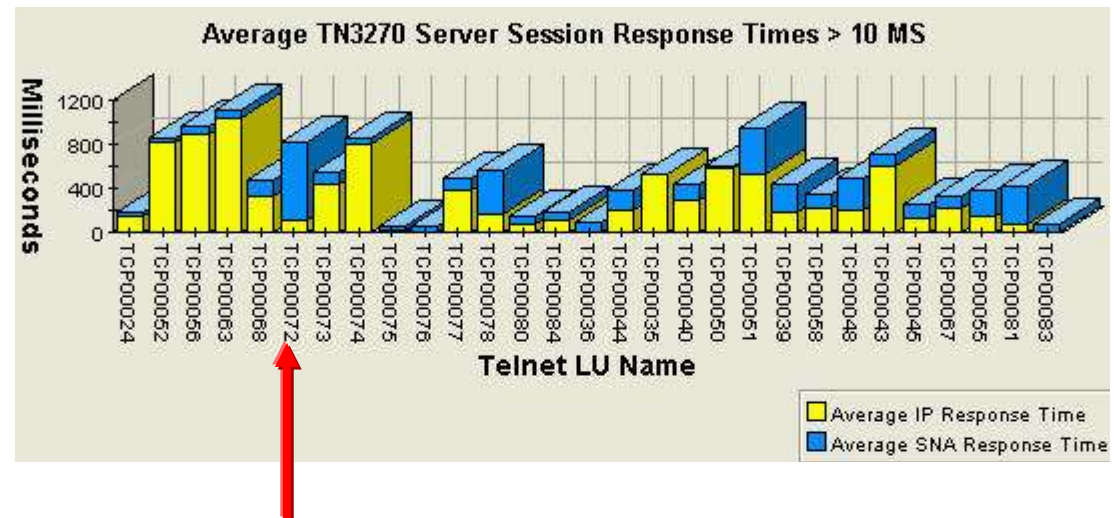


Erratic TN3270 response times ...

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

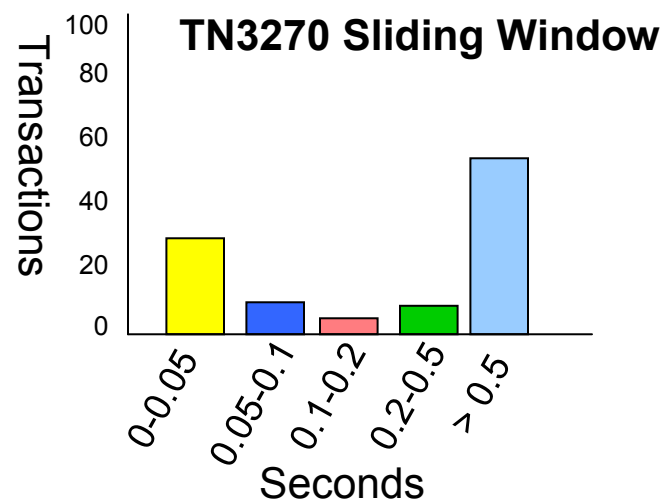


2. The average response time and average SNA response time are fairly high. In contrast, average IP response time is good, so does not appear to be a network problem.



Erratic TN3270 response times ...

- 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.
- 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.

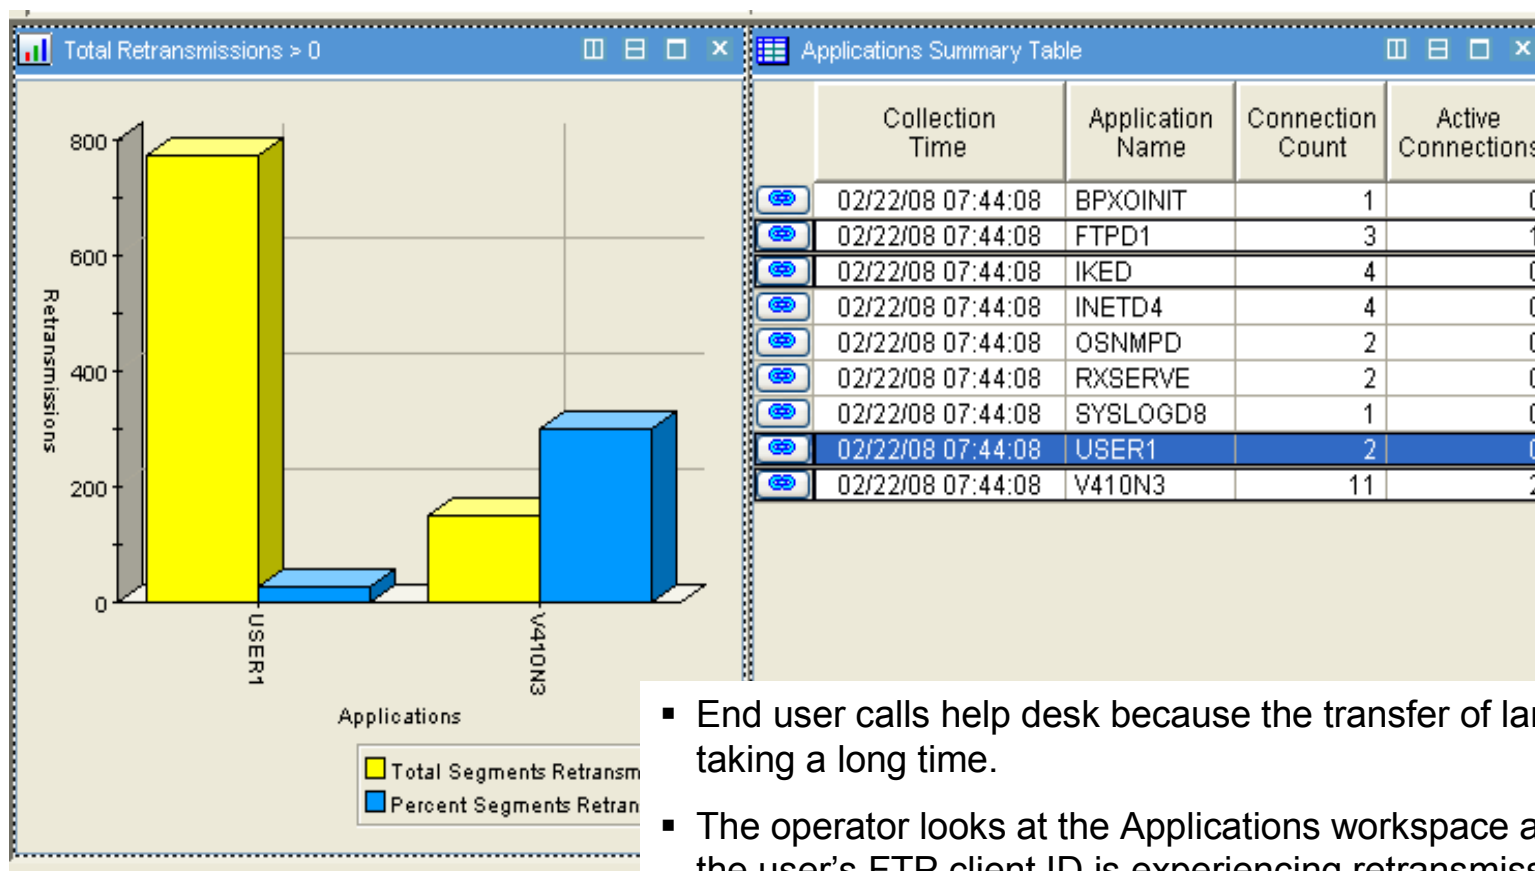
z/OS IP Security Support



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

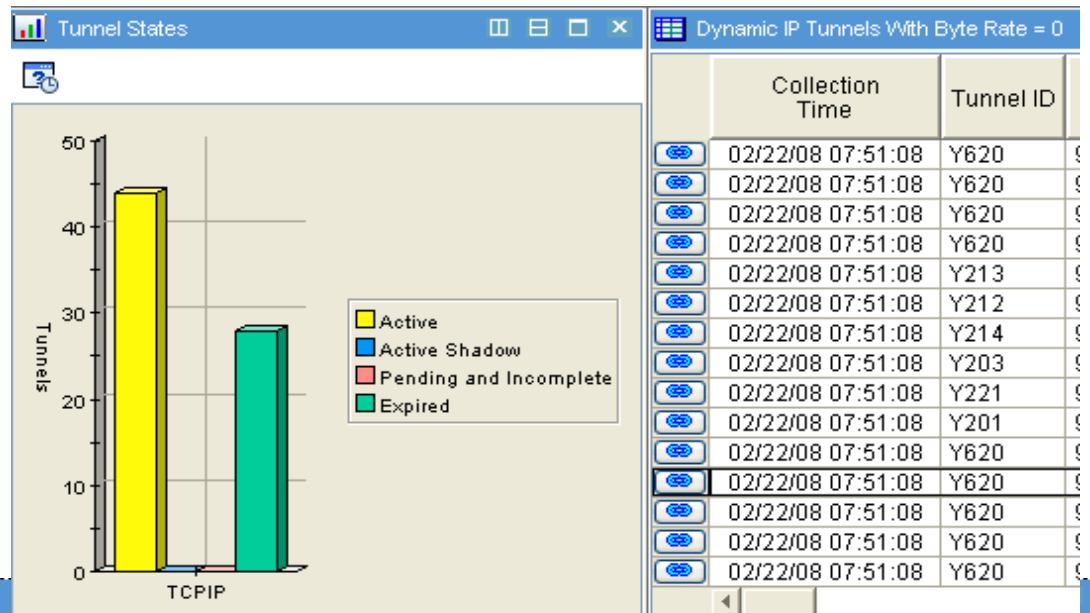
Application performance problem ...



- End user calls help desk because the transfer of large files is taking a long time.
- The operator looks at the Applications workspace and sees that the user's FTP client ID is experiencing retransmissions.

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.



Tunnel ID	Local Security Endpoint	Remote Security Endpoint	Total Inbound Packets (in G)	Total Inbound Packets	Total Outbound Packets (in G)	Total Outbound Packets	Total Packets (in G)	Total Packets	Inbound Packets	Outbound Packets	Packets	Packet Rate
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	
Y620	9.42.45.206	9.42.45.114	0	0	0	0	0	0	0	0	0	

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

Dean Butler (butlerde@us.ibm.com)

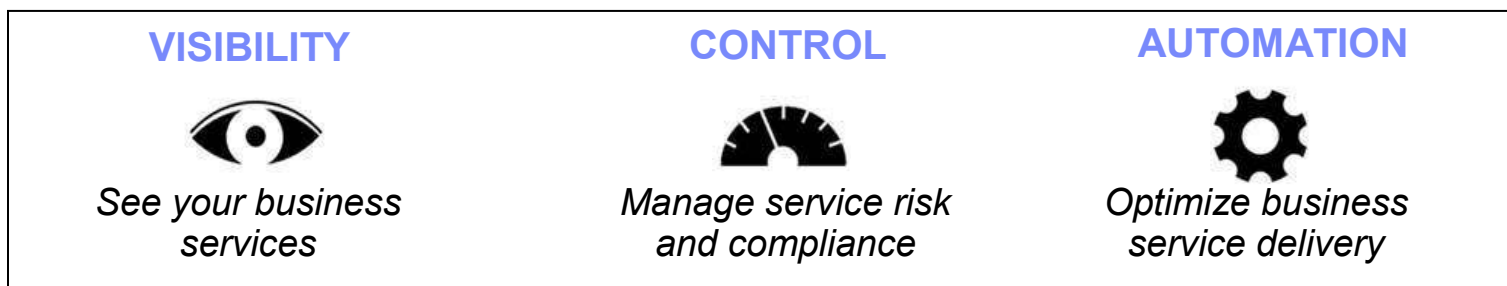
IBM's Integrated Service Management approach is recognized as best in class

Integrated Service Management



IDC Market Share rankings:

- #1 Overall in Systems / Network Management
- #1 in Overall Performance and Availability Mgt.
- #1 Performance Management
- #1 Event Automation
- #1 Network Management
- #1 Output Management
- #1 Archiving
- #1 Identity and Access Management
- #1 Security and Vulnerability Management
- #1 Enterprise Asset Management



Learn more - <http://www.ibm.com/software/tivoli/features/zsmc/>