This paper was originally published in the Proceedings of the Computer Measurement Group's 2005 International Conference.

DDF Performance Analysis – Does it really have to be this complicated?

Bob Chaney
Delta Technology, Inc.
Bob.chaney@delta.com

DB2 DDF processing, where the database is centralized and the applications are distributed, can offer some interesting challenges when diagnosing "End-To-End" performance problems. This paper uses a number of different sources from SMF to UNIX log files to web reports, to create a single view of a distributed workload as part of a performance problem diagnosis..
Warning: Success is NOT guaranteed!

Introduction

The biggest I/T initiative at Delta Technology over the past 3 years has been the creation of an Enterprise Data Store (EDS). This large DB2 environment contains different business subject areas and is designed to provide "near-real-time" business information. The architecture is a centralized database with many distributed application servers, tied together with a gigabit Ethernet network. The software conduit is DB2 Connect on the UNIX and/or Windows application servers and the Distributed Data Facility (DDF) on the zSeries database servers. The other software deployed in the environment is Tuxedo, which is a teleprocessing (TP) monitor, Crystal Reports on Windows and MQ series.

The EDS environment cuts across several infrastructure and application support silos in the organization, requiring lots of "East-West" communications, especially when diagnosing end-to-end performance problems. In our case the responsibility is distributed across the following departments:

Support Department	***Infrastructure Responsibility***
Large Systems Engineering	*zSeries hardware, Capacity & Performance, DB2 Systems Programming*
Mid-Tier Engineering	*UNIX hardware, Capacity & Performance*
Workgroup Engineering	*Windows hardware, Capacity & Performance*
Network Engineering	*Network hardware, Capacity & Performance*
Middleware	*MQ Series, Tuxedo*
Data Services	*DB2 database design/maintenance, DB2 Connect*
Systems Operations	*Incident management, Problem management*

The responsibility of Systems Operations, under our partial Information Technology Infrastructure Library (ITIL) deployment, is Incident, Problem and Change Management. This means they coordinate the "East-West" communication of an incident, or problem as well as manage the acceptable to work a performance problem. The rest of this paper will describe an actual performance problem, the process by which it was diagnosed using information from several silos, and finally, show how ITIL Configuration Management might have reduced the diagnosis and root cause discovery time.

The Application

The application is designed to process passenger electronic ticket (ET) data. Ticket information is first captured in the Transaction Processing Facility (TPF). From there it is transported to a mid-tier application server using MQ Series. The

.

I/T Change process. Even with this process deployed, however, it still takes longer than application server processes the ticket and either inserts/updates/deletes the ticket in the EDS. The mid-tier application is designed to handle 20 concurrent paths of ticket data and is instrumented to track the number of tickets queued in each path. The queue length of the paths is the primary early warning metric for performance problems. In addition to transaction data flows there is a daily (4 AM) TPF file maintenance process that "dumps" ticket information into the pipeline.

The Physical Architecture

Figure 1 shows the physical architecture on which the ET application resides

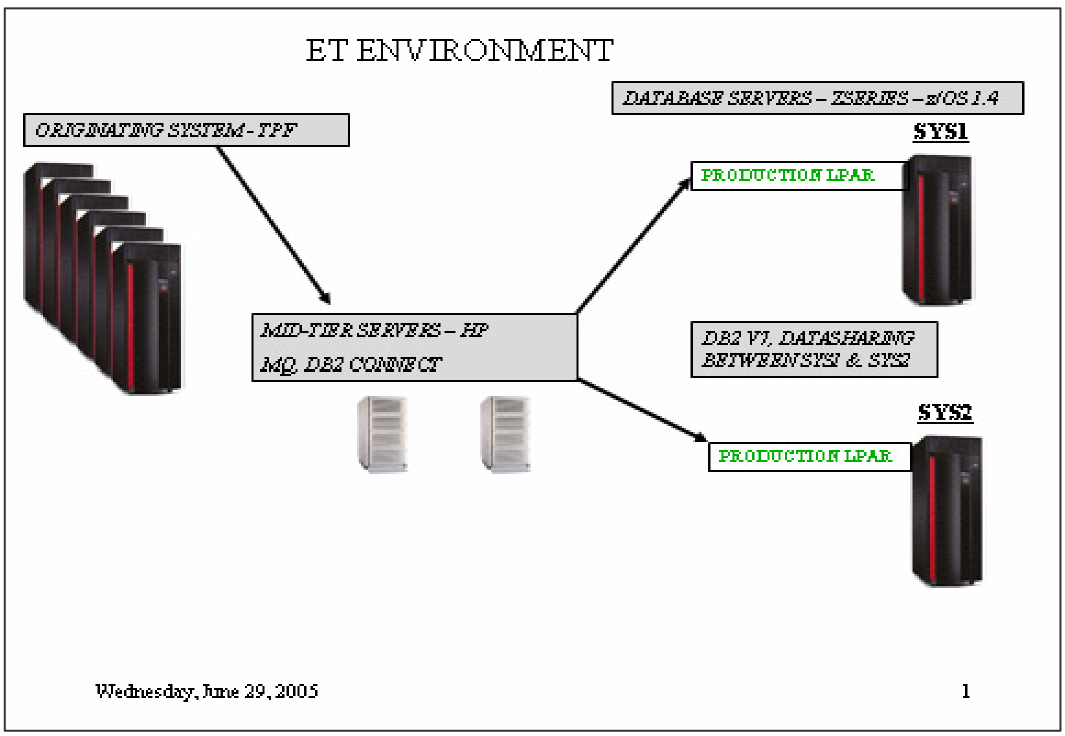

Figure 1. Physical Architecture

This physical architecture is viewed logically as a pipeline. Data flows from TPF to the Mid-tier servers where it is processed 20 paths at a time and inserted/updated/deleted in the DB2 database. The ET transactions are identified in DDF by their unique DB2 *authid*, which allows us to monitor and manage them in specific Workload Manager (WLM) service and report classes. The entire environment is connected via a gigabit Ethernet network.

The Problem and Diagnosis

The ET performance problem was first noticed on Sunday, April 3, 2005. Monday and Tuesday the problem was investigated by the application team and data services with no resolution. The problem was first made known to the Large Systems group (the author's department) on Wednesday, April 6 via email from Data Services. The email reads:

> *Bob, ET has experienced sporadic slowdowns since Sunday. I ran some DB2 Connect traces and reported the problem to IBM. The errors we are seeing on the DB2 Connect side indicate that the network timeout limit has been reached. IBM has suggested that we look at WLM and see if there are any errors or indications of slow performance from your side. Will you check into this when you get a chance and let me know what you see?*

This email prompted several phone calls and additional emails to discover that the Mid-Tier folks had taken a look at their servers and found nothing amiss, Data Services DBAs had looked at the database and found nothing noteworthy. In addition, all changes for the prior weekend had been examined, and the only change that should have affected **this environment** (*emphasis added – we'll revisit this later in the paper*) was the "spring forward" time change on the morning of April 3rd. So now it is my turn to take a look. The first metric we want to review is ET response time. Figure 2 shows the hourly response time for the previous week and the current week-to-date.

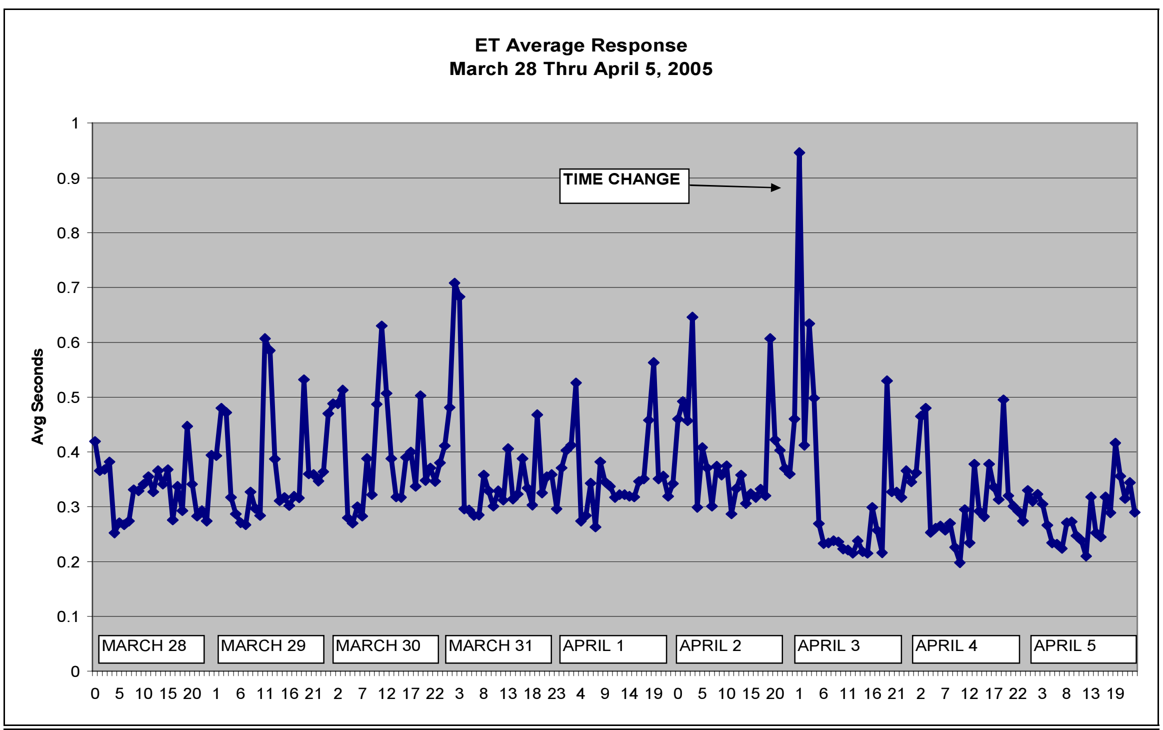

Figure 2. Average Response

The average response time actually looks better this week. The only glaring exception occurred during the morning of April 3rd, and we know that was the result of stopping the data flow while all the system clocks were moved forward one hour. While working with the Type 72 goal records I did notice some ET Performance Indices (PIs) greater than 1.0, indicating they were occasionally missing their .5 second response target. This is also obvious from the chart. As a result, I made some minor changes to the ET Service Class definition in WLM, giving them a higher importance. However, I did not expect this was the root cause of the ET performance problem. Subsequent emails on April 7th confirmed this and all the groups who have investigated their responsible areas agreed that it was now time to involve Network Engineering. This leads us to a discussion about the performance problem diagnostic process.

The Process

Our diagnostic process can be described as the "Not me, not me, not me…… let's look at the network" process. Since all of the server environments are well instrumented with RMF/SMF/MXG on the database servers and HP MEASUREWARE/SERVICE REPORTER on the application servers, performance data is readily at hand for diagnosing performance problems. The network, however, uses bandwidth as the primary performance metric and tends to require sniffer traces for more useable performance diagnostic information. For that reason we usually wait until we've exhausted all the "usual suspects" like CPU utilization, disk bottlenecks, memory constraints before we involve the network folks. As we'll see later in the paper, this was a mistake.

Problem Diagnosis, Part Two

Network Engineering was invited to the party via the following email:

We're currently experiencing performance problems on ET, and we'd like to get someone from the Network group involved to help us troubleshoot the problem. Here is some background information and more details about the problem we're seeing:

- *ET started experiencing slow response on Sunday, 4/3 – sporadic slow response has continued since then.*
- *The applications use DB2 Connect to a DB2 database on z/OS.*
- *DB2 Connect error log is filled with TCP/IP connect errors – DIA3202C the TCP/IP call "connect" returned the errno="238" – the 238 indicates that the connection has reached the network timeout limit and is terminated by the network.*
- *No changes (that we are aware of) have been made on the application server, DB2 Connect, application code, z/OS or DB2.*

By now it is Thursday, April 7th and the problem, while not critical, is causing our "near zero latency" requirement for electronic tickets to be dangerously close to being a missed target. It takes some time for the sniffers to be deployed because they must be located at very specific points in the network, and only packets between the two interesting nodes will be captured. In our case it was the application server, several

Network switches and the two database servers (SYS1 & SYS2) that were to be "sniffed". The sniffers were in place and capturing data by the morning of April 8th. Analysis of the packets caused the investigation to move in a completely different direction that led us to the root cause very quickly.

The Root Cause

Sniffer data allows the analyst to open the data packet and peer inside. You see routing information and, in our case, the actual distributed SQL statements that are being sent. In the case of the ET performance issue the sniffer data uncovered what looked to be strange routing information from one of the network sub-net switches. This information prompted another look at the changes that occurred during the previous weekend, and brought to light a seemingly unrelated change to a completely different sysplex. That change implemented Virtual IP Addressing (VIPA) in a sysplex that shares nothing (or so we thought) with the sysplex in which the EDS database runs. Figure 3 shows both the EDS database servers along with the separate PRODPLEX servers. As we work through the root cause analysis we'll add information that will highlight the ultimate root cause.

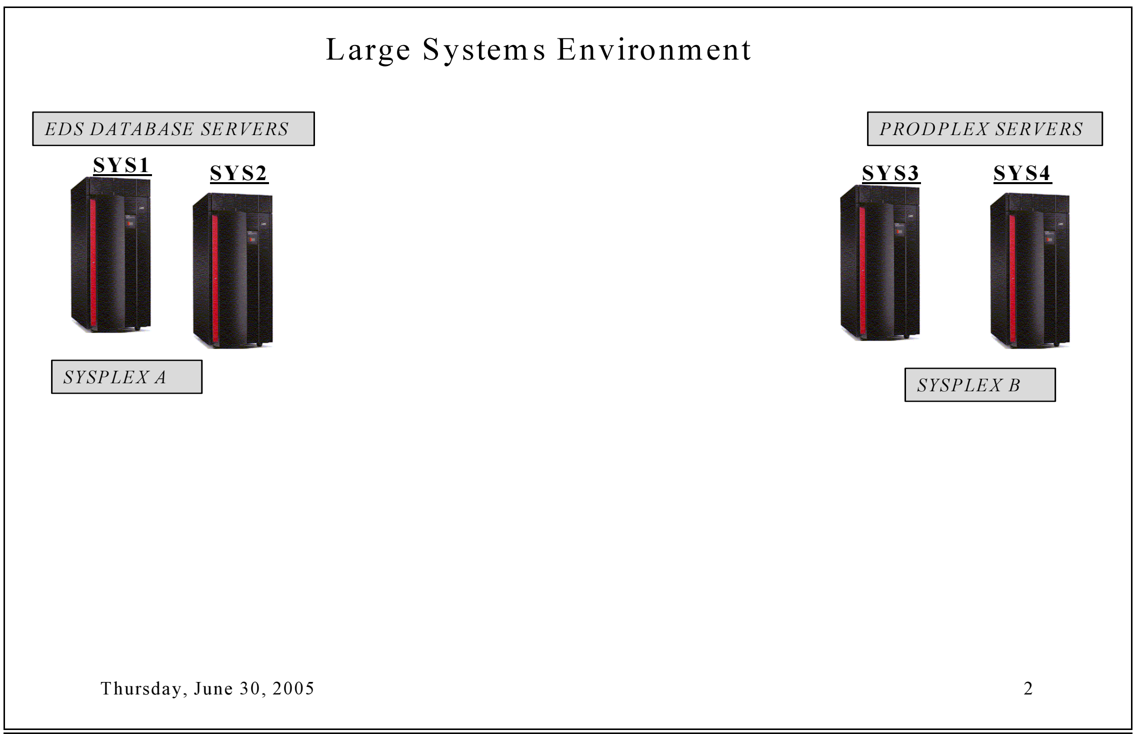

Figure 3. EDS and PRODPLEX

VIPA provides fault tolerance by using a virtual device and a virtual IP address in the z/OS TCP/IP stack rather than linking a physical device/IP address. The virtual device is always active and never sees a failure, so packets will be sent/received successfully as long as one of the OSA adapters on the zSeries machine is working. This prevents packet failures even if an adapter fails. VIPA requires a change in network protocol from static route to OSPF (Open Shortest Path First) routing. While VIPA is implemented at the sysplex/LPAR level, OSPF is implemented at the network sub-net level, and that is where the root cause of the ET performance problem revealed itself. Figure 4 adds the network components and ET application server to the diagram.

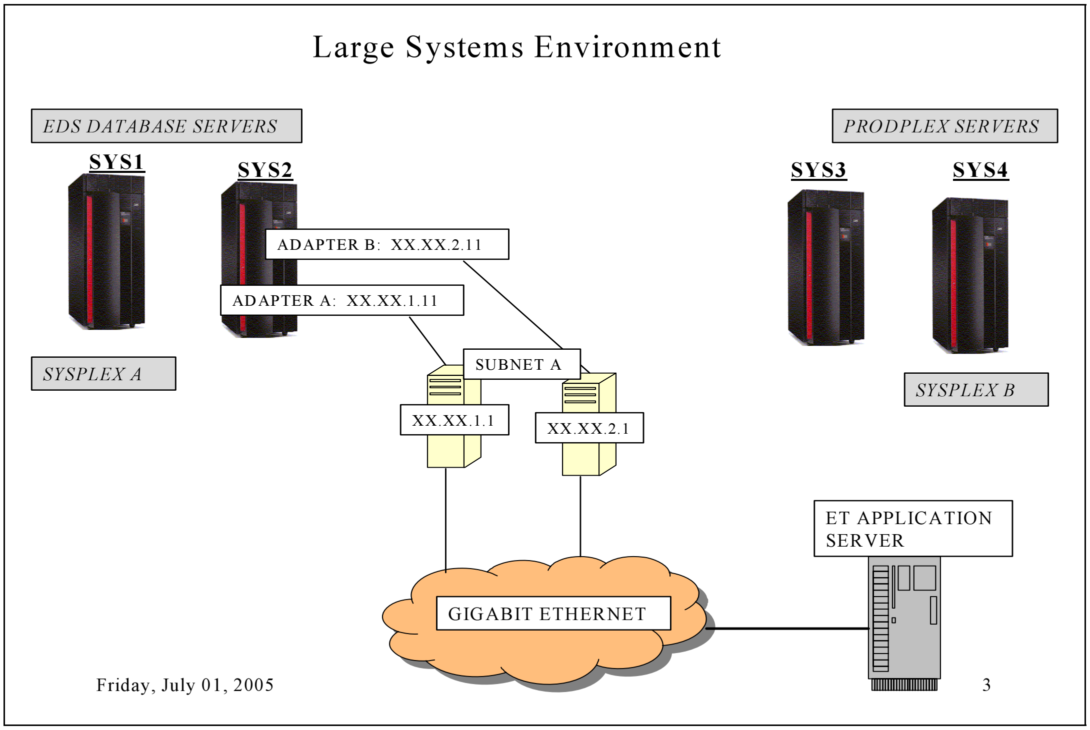

Figure 4. ET Network and Server Diagram

From the diagram in figure 4 we see that subnet A consists of two switches with TCP/IP addresses of xx.xx.1.1 and xx.xx.2.1. This is used by transactions going from the ET application server to the DB2 datasharing member on SYS2. A separate sub-net (not shown) is used to route transactions going to SYS1. Information provided to DB2 Connect by WLM in Sysplex A determines which datasharing member the transaction will use to access the database.

What was revealed at this point (Friday afternoon) is that the VIPA change to Sysplex B caused a network change to the 1.1 / 2.1 subnet, implementing OSPF. Figure 5 adds the rest of the network information and gets us closer to the root cause. The implementation of an OSPF area along with the network connections of this subnet to a zSeries machine in Sysplex B (which was there LONG before the OSPF change, but not realized), now gives us the answer to the “what changed?” question. The implementation of VIPA on Sysplex B and the required OSPF change to the subnet is now the most likely root cause candidate.

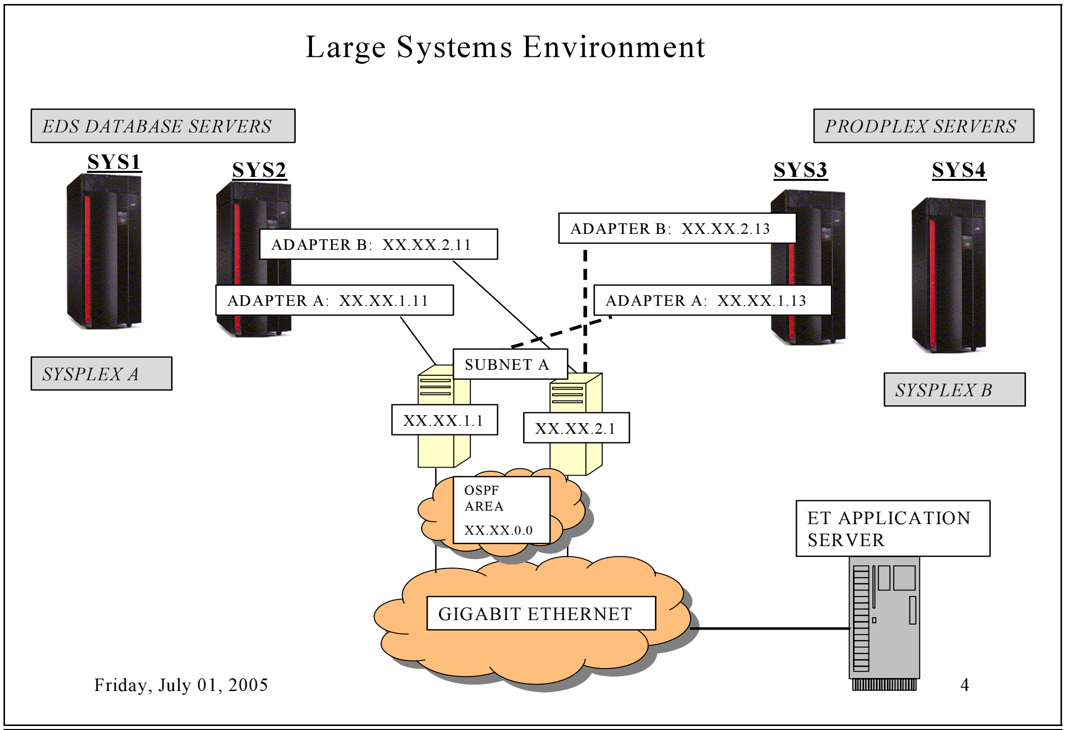

Figure 5. OSPF Area and Sysplex B Connections

Referring to Figure 5, the ET transactions will flow through either/both of the two network switches in subnet A. What the sniffer data revealed was that some ET transactions using the xx.xx.1.1 switch were actually being routed to SYS3 in Sysplex B, and of course they were failing (hence the network timeouts). Apparently, the OSPF routing logic thought this was the shortest path for those transactions, even though the DB2 system they needed was in Sysplex A. The reason for this mid-direction in routing has yet to be determined. The solution, however, was simple.

The Solution

Network Engineering determined, through testing, that if a layer 2 connection was installed between the "1.1" and "1.2" switch in subnet A that the previously misdirected transactions would be routed to xx.xx.1.1 to xx.xx.2.1 to SYS2, where they would complete successfully. The cable was installed on Saturday morning and the ET performance problem was solved. While this was good news, the fact that it took 6 days was not acceptable, and we believe the way to reduce that time is to implement an additional ITIL processes.

The ITIL Connection

You will find a number of papers about ITIL (Information Technology Infrastructure Library) in this year's CMG conference proceedings, so I won't go into much detail. ITIL provides a process framework by which IT infrastructure service delivery can be improved. The processes are divided into two core areas with five sub-areas in each. Figure 6 below shows the ITIL core areas and sub-areas.

SERVICE SUPPORT	SERVICE DELIVERY
SERVICE DESK (FUNCTION)	
INCIDENT MANAGEMENT	SERVICE LEVEL MANAGEMENT
PROBLEM MANAGEMENT	FINANCIAL MANAGEMENT
CHANGE MANAGEMENT	CAPACITY MANAGEMENT
CONFIGURATION MANAGEMENT	I/T SERVICE CONTINUITY MANAGEMENT
RELEASE MANAGEMENT	AVAILABILITY MANAGEMENT

Figure 6. ITIL Core and Sub-Core Areas

The ITIL process we are most interested in as a means of reducing performance problem diagnosis time is Configuration Management. As previously mentioned, we have implemented Incident, Problem and Change Management, but for this particular problem Configuration Management would have been very helpful.

Configuration Management processes revolve around the existence of a Configuration Management Data Base (CMDB). Configuration Management will:

- Account for all IT assets
- Provide accurate information to support other Service Management processes
- Provide a sound basis for Incident, Problem, Change and Release Management
- Verify infrastructure records in the CMDB and correct exceptions.

The CMDB consists of Configuration Items(CIs) that include not only the identity of IT assets but also their relationship to one another as they are deployed to provide IT service to the customer. In our case that customer would be the ET application. The CIs would be the servers and network nodes that are used by the ET customer. Some of these same CIs would also apply to customers of SYS3 in Sysplex B.

If a CMDB were used as part of the Change Management process it is possible the sharing of subnet A between the two sysplexes would have been noted during the weekly Change Action Board meeting, and the potential impact to the EDS environment would have been listed as a known risk. At the very least, a CMDB could have been used in the early stages of Problem Management and we might have discovered the subnet issue more quickly.

In June, 2005 we deployed our first CMDB and we are actively using and improving it.

Conclusion

The new distributed application, centralized database environment of our Enterprise Data Stores has introduced a new level of complexity that previously did not exist in our I/T infrastructure. We can no longer diagnose performance problems within each infrastructure's silo and hope we spot the problem, but instead need to have a process and tool that identifies all the infrastructure components and their relationships. As the incident in this paper points out, the old "not me, not me…..must be in the network" process takes too long, especially if the root cause is in the network.
In the future we hope to use the CMDB as a diagnostic tool whenever we engage in the Incident, Problem and Change Management processes.

References

Leveraging z/OS TCP/IP Dynamic VIPAs and Sysplex Distributor for higher availability, IBM White paper

Dynamic VIPA Routing & Workload Balancing in a DB2 Datasharing Environment, IBM White paper – www.ibm.com/software/data/db2/os390/support.html

Information Technology Infrastructure Library, ITIL institute CD

Thanks to Betsy Simons for DB2 Connect advice

Thanks to Matt Grimes for VIPA, OSPF and general network advice.