

Staying Ahead of Network Problems at DTCC

Nalini Elkins
Inside Products, Inc.

Richard Warren
Depository Trust & Clearing Corporation

Wednesday, August 10, 2011
Session Number 9289

Our SHARE Sessions – Orlando



- 9285: TCP/IP Performance Management for Dummies
Monday, August 8, 2011: 11:00 AM-12:00 PM
- 9269: IPv6 Addressing
Wednesday, August 10, 2011: 11:00 AM-12:00 PM
- 9289: Staying Ahead of Network Problems at DTCC
Wednesday, August 10, 2011: 3:00 PM-4:00 PM



Agenda



- Introduction to DTCC
- Business requirements
- Workload monitoring
- Case study
- Proactive management
- Processing of alerts and warnings

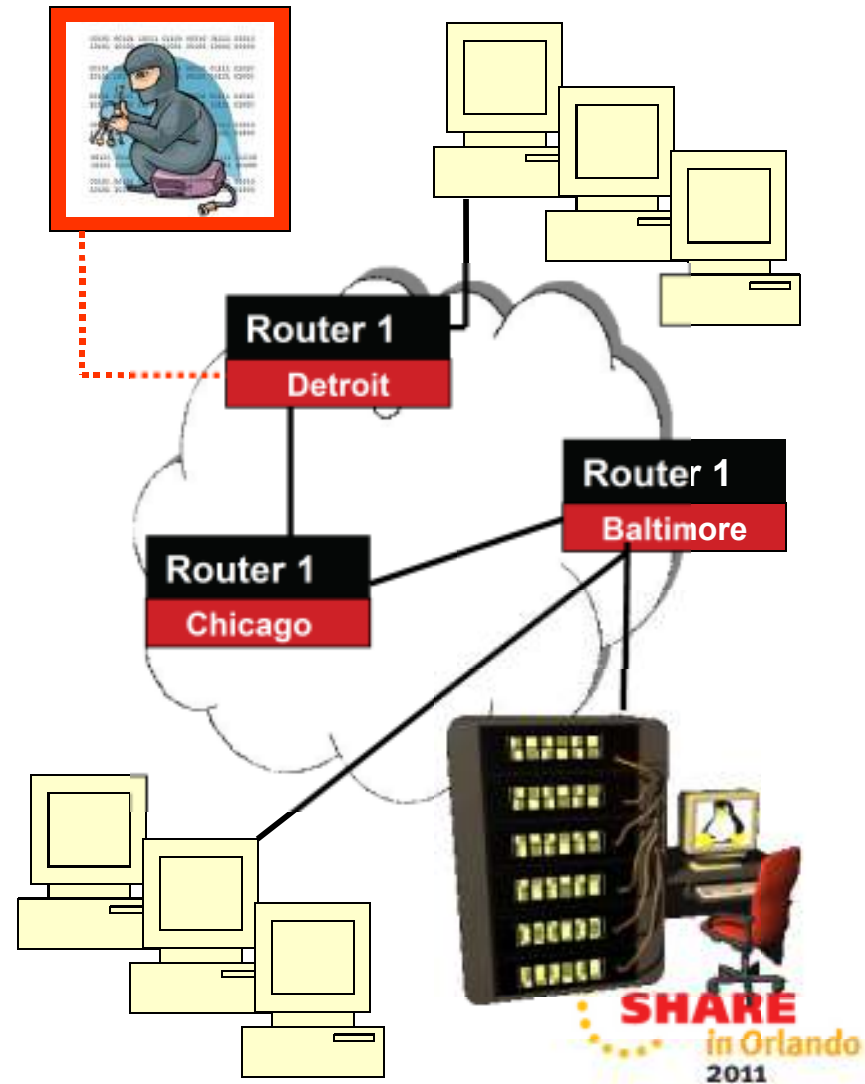


Introduction to the DTCC


- The Depository Trust & Clearing Corporation (DTCC) is at the epicenter of the financial world.
- The business of DTCC involves the safe transfer of securities ownership and settlement of trillions of dollars in trade obligations, under tight deadlines every day.
- At the same time, DTCC's primary mission is to protect and mitigate risk for its members. DTCC ensures the capacity, certainty and reliability required to clear and settle today's enormous trading volumes.

Business Requirements

- Interconnect the financial world...
- We are a service provider
- Close the markets...
- Do all this in a timely manner
- And... run it as a business
- Let's take each of the above and see the implications for network management.

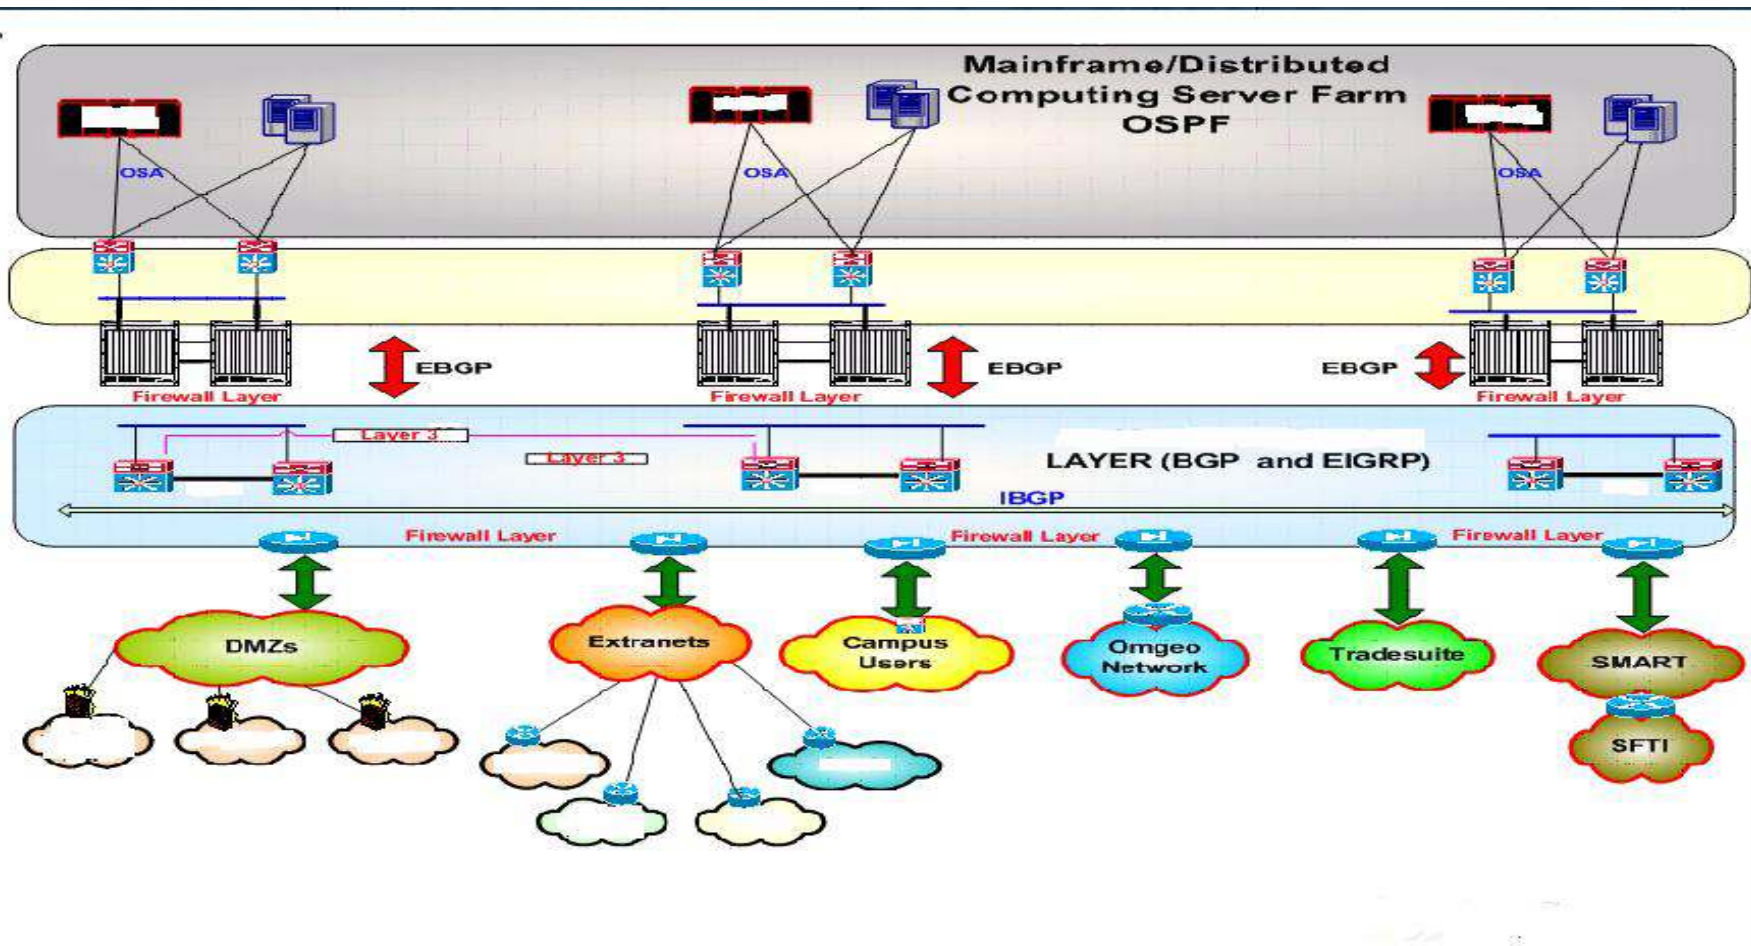


DTCC Interconnects the Financial World



- The network is at the heart of DTCC's business.

High Level Network Diagram



DTCC is a Service Provider

- What kind of service are we providing?
 - View by business partner
 - View by service (port)
 - Monitor availability
 - Monitor network response time
- How do we know if we are providing it?
 - Set thresholds
 - Define services
 - Get alerts
 - Monitor
- What do we do if there is a problem?


Case Study

- The client is complaining of much longer times to complete their work.
- Doing mostly NDM transfers
- Work is highly critical
- Money is at stake



How do we solve this?

- What data do we need?
 - By IP address
 - By local port
- Where do we get it from?
 - SMF 119
 - MICS / MXG
 - Inside the Stack
- Keeping and storing data is a large task



We use this data daily!

- We create daily reports for:
 - Retransmits
 - Duplicate acks
 - Session duration
 - High bytes in / out

ITS: TOP 50 DUPACKS



START TIME	REMOTE IP ADDRESS	REMOTE PORT	LOCAL PORT	BYTES IN	BYTES OUT	SOCKET NAME	SESSION DURATION	ROUND TRIP TIME	ROUND TRIP VARIAN	RETRAN COUNT	DUP ACKS	
5:17:00	1.2.3.	089	1414	12399	2385K	605M	MQMPCHIN	15:48:05	32	28	173	251459
3:44:10	1.2.3.	100	1422	59279	1181K	893M	MQMPCHIN	17:32:37	27	3	168	229039
3:25:46	1.2.3.	096	1414	55650	234K	305M	MQMPCHIN	16:37:38	68	118	117	170095
5:52:22	1.2.3.	36	1614	19169	621K	255M	MQMPCHIN	14:08:59	82	53	3	131342
5:52:44	1.2.3.	47	1416	19265	1102K	284M	MQMPCHIN	14:07:42	108	88	60	116438

Benefits of storing in DB2

- We are near 700 MILLION detail records
- Scalability of solution and integration of data took quite a while to achieve
- Used DB2 partitioning and indexing to improve performance
- Can access the data via SAS or SQL (JDBC)
- Use the data for trending as well as diagnostics

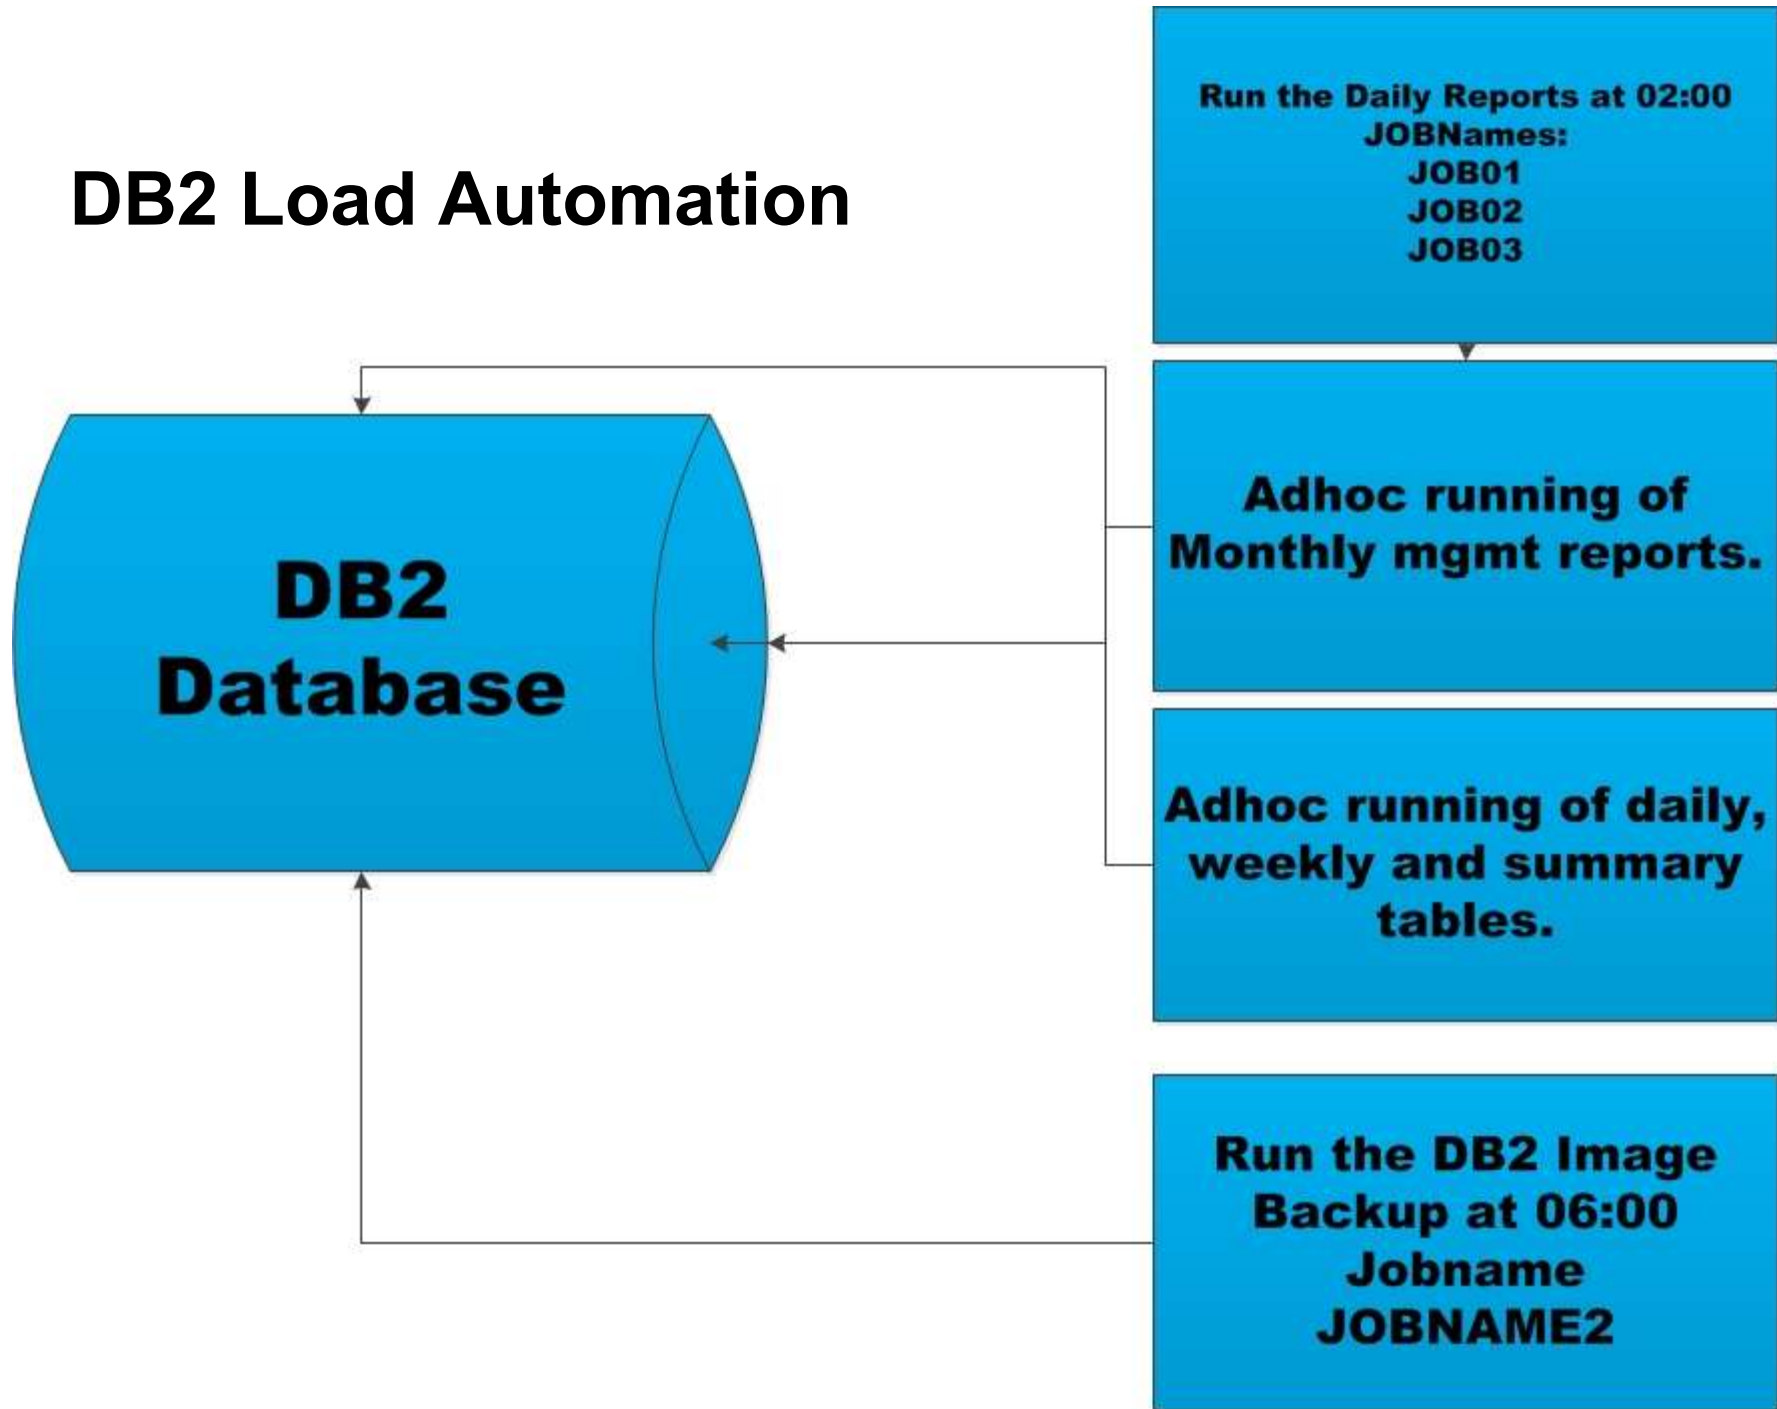


Capture the Workload Data

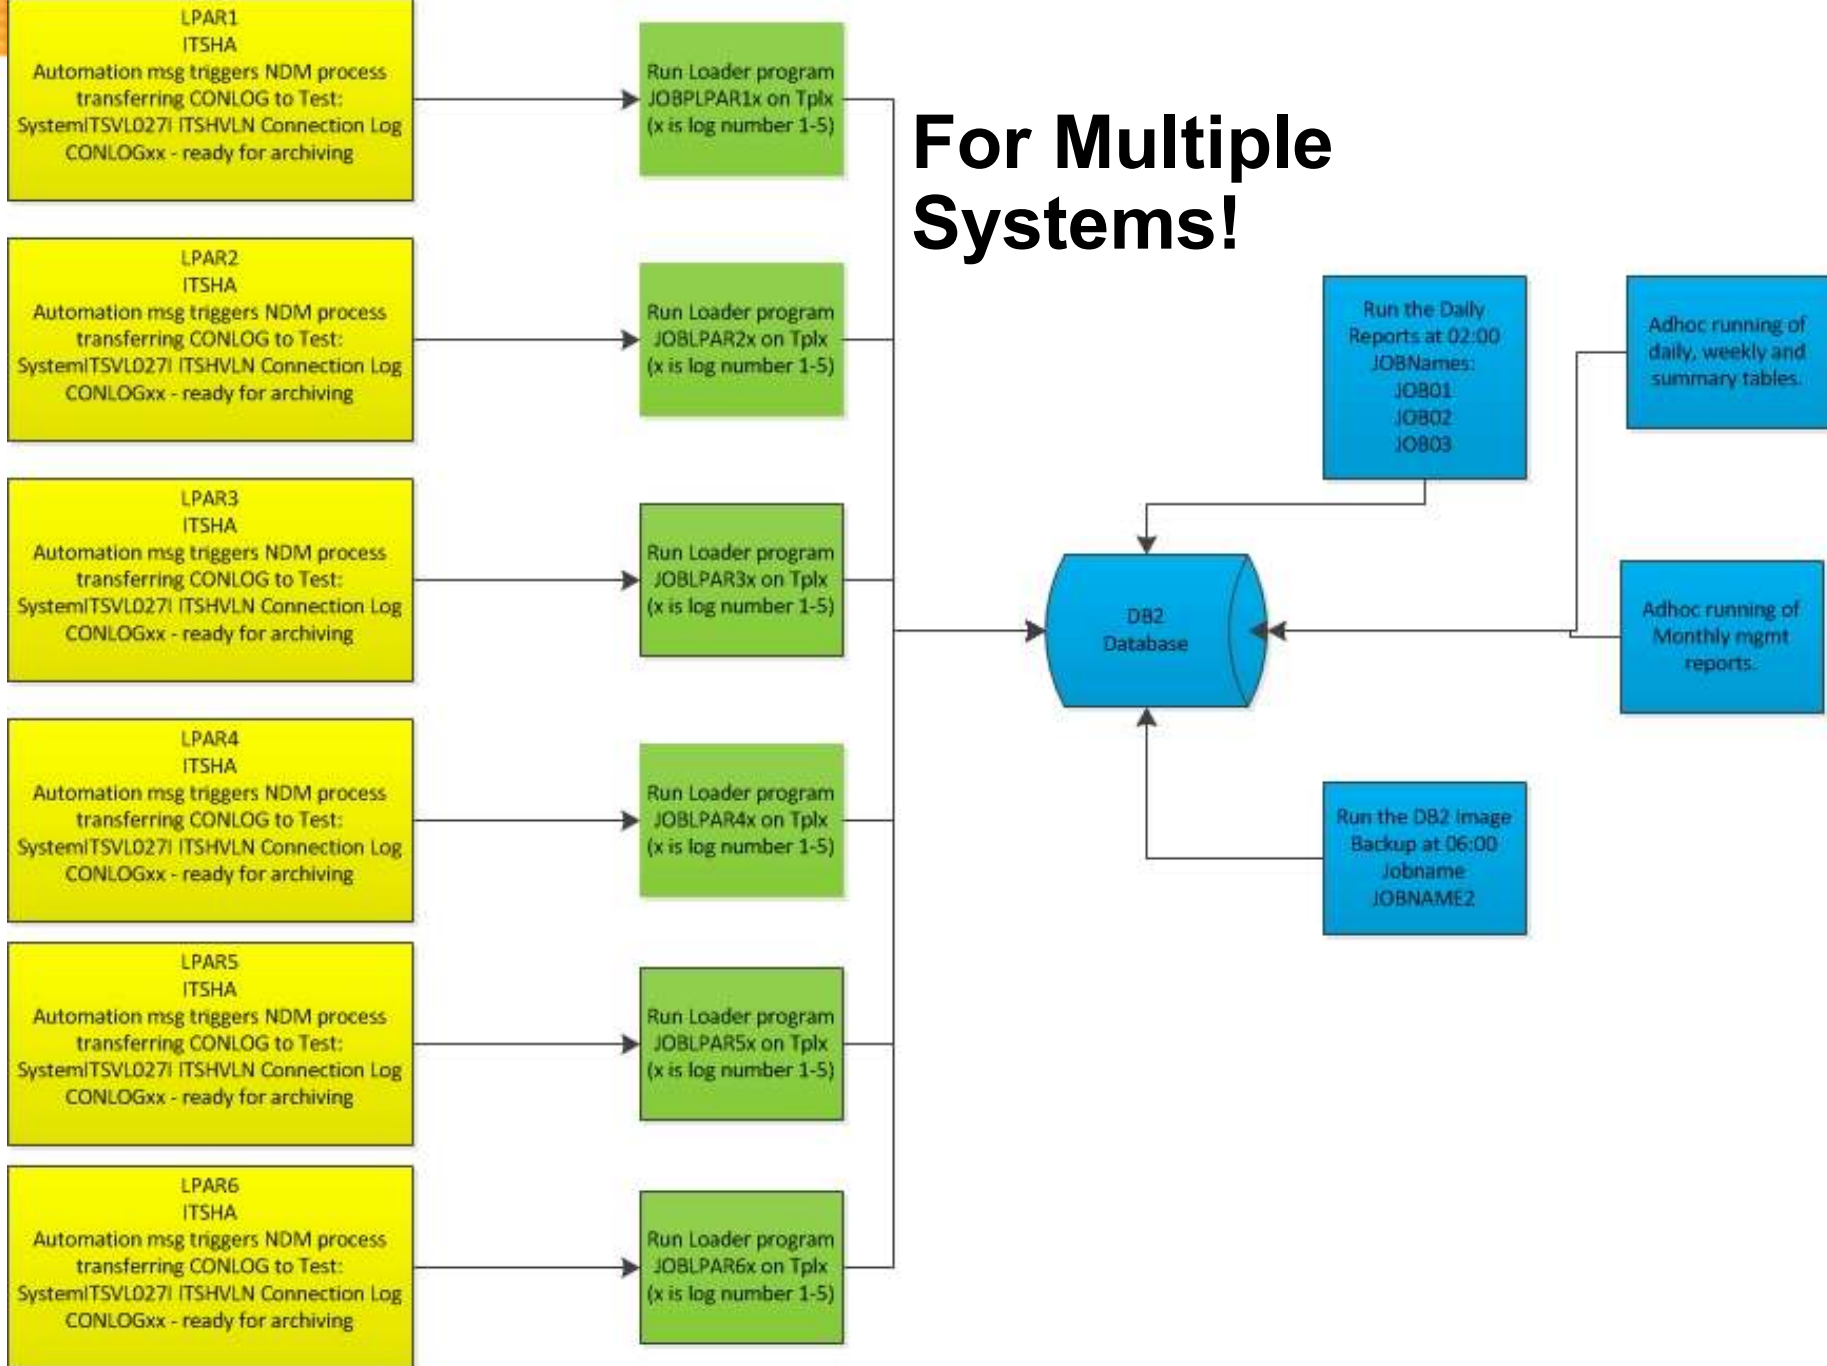
**LPARx
ITSHA
Automation msg triggers
NDM process
transferring CONLOG to
Test:System
ITSVL027I ITSHVLN
Connection Log
CONLOGxx - ready for
archiving**

- Started task ITSHA (from Inside Products) captures TCP workload data from IBM Network Management API
- Data is stored in VSAM linear (memory mapped VSAM). This is the Connection Log for the day.
- At the end of the day, the Connection Log is loaded into DB2.

DB2 Load Automation



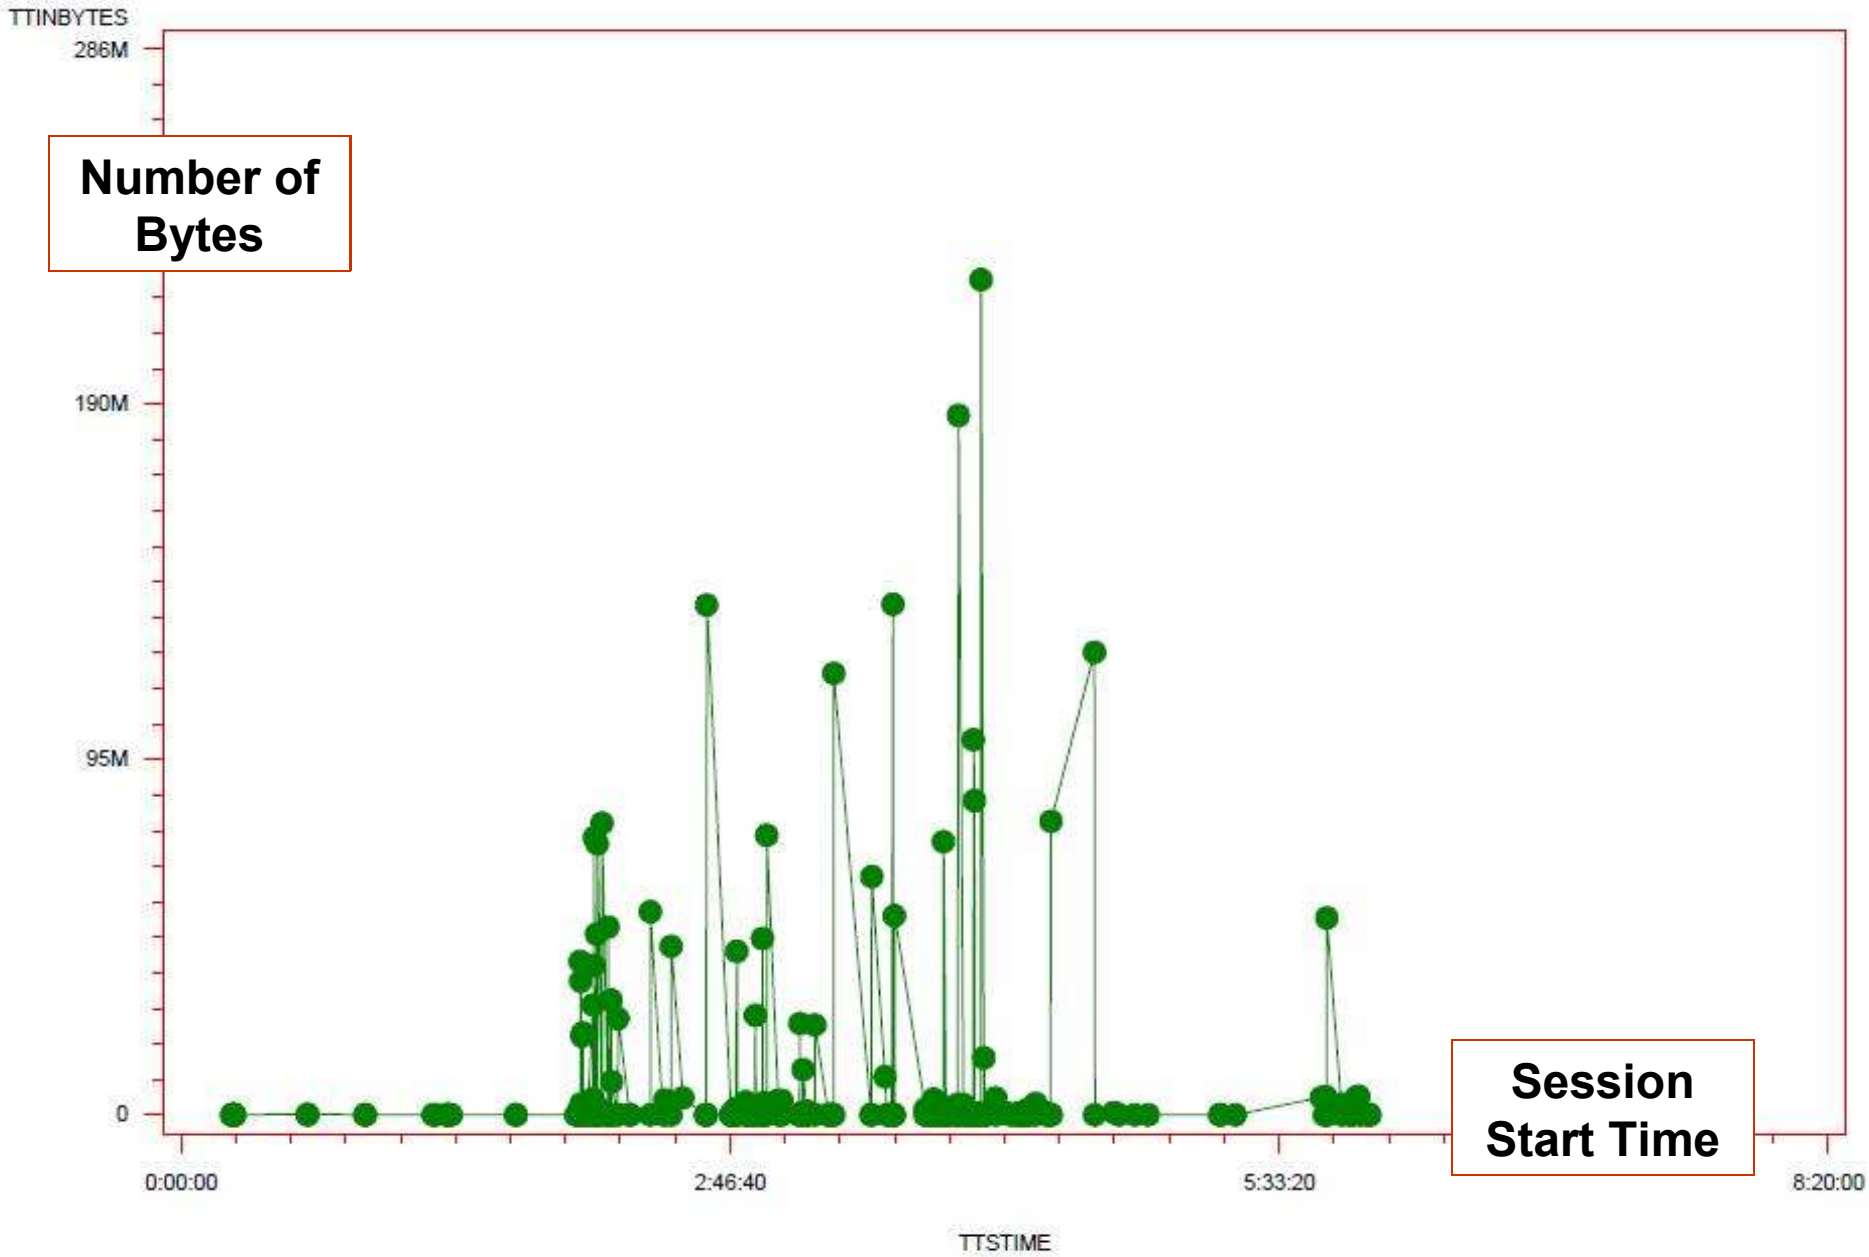
For Multiple Systems!



SAS Reports from DB2

- Now, we have a database.
- Still need reporting capability.
- We created our own SAS reports.





We solved the problem...

- Client had slow response problem
- It was not a bandwidth Issue
- The issue was an active SSL trace in the NDM
- But, we needed a better reporting system



The Challenge with SAS Reports


- Writing SAS graphics code is time consuming and cumbersome.
- Is this what a network analyst should be doing?
- Every change requested (time, graph type, etc.) is a code change!



New Graphical Reporting System



Session Start Time: 08:00:00 - 14:00:00
Session End Time: 00:00:00 - 23:59:59
Remote Host: 1.2.3.4
Sorted By: Start Time Descending



Benefits of New GUI

- Retrieving data via the interface is 2-3 seconds
- Learned about partitioning and indexing to improve performance
- The DB2 interface creates various graphs (bar, area, 3D etc) as well as text

Submit

Stack ID: <input type="text"/>	Local IP: <input type="text"/>
Local Port: <input type="text"/>	Remote Port: <input type="text"/>
Socket Name: <input type="text"/>	Remote IP: <input type="text"/>

Entries to Show: Sort By: Sort Order:

Session Start

Date: <input type="text"/>	Start Time: <input type="text"/>	End Time: <input type="text"/>
Date Format: YYYY-MM-DD	Time Format: HH:MM:SS	

Session End

Date: <input type="text"/>	Start Time: <input type="text"/>	End Time: <input type="text"/>
Date Format: YYYY-MM-DD	Time Format: HH:MM:SS	

Graph Options

Graph Type: <input type="text" value="Area"/>	3D: <input type="text" value="Yes"/>
<input checked="" type="radio"/> Create Multiple Graphs	<input type="radio"/> Create Single Graph
Height: <input type="text" value="300"/>	Width: <input type="text" value="850"/>

Variables to Graph

<input checked="" type="checkbox"/> BytesIn	<input type="checkbox"/> BytesOut
<input type="checkbox"/> Retransmits	<input type="checkbox"/> Duplicate Acks
<input type="checkbox"/> Round Trip Time	<input type="checkbox"/> Round Trip Variance
<input type="checkbox"/> Session Duration	

Graph Labels

Business Entity Correlation

- DTCC has many partners.
- So, we added a new table to the DB2 database.
- When diagnosing a problem or trending, it is very useful to see the business partner name as well as the IP address.



Reports by Business Entity

Details - TCP Connections
 Stack ID: P001
 Session Start Date: 2011-06-15
 Session Start Time: 00:01:15 - 23:59:59
 Session End Time: 00:00:00 - 23:59:59
 Sorted By: Start Time Descending



	Business Entity	Start Time	End Time	Stack	Local Host	Local Port	Socket Name	Remote Host	Remote Port	Bytes In
1	Broker # 1	2011-06-15 00:01:15	2011-06-15 00:01:17	P001	192.168. []	1364	NDMCF2	1.1.1.1	49814	4,799
2	Bank # 1	2011-06-15 00:01:15	2011-06-15 00:01:17	P001	192.168. []	63743	NDMCF2	2.2.2.2	1376	3,372
3	Broker # 2	2011-06-15 00:01:17	2011-06-15 00:01:18	P001	192.168. []	1364	NDMCF2	1.1.1.1	49816	4,799
4	Exchange #1	2011-06-15 00:01:17	2011-06-15 00:01:19	P001	192.168. []	63749	NDMCF2	2.2.2.2	1376	3,372
5	Broker # 1	2011-06-15 00:01:17	2011-06-15 00:01:19	P001	192.168. []	63748	NDMCF2	1.1.1.1	1364	11,262
6	Bank # 3	2011-06-15 00:01:17	2011-06-15 00:01:31	P001	192.168. []	1364	NDMCF2	2.2.2.2	52230	4,074
7	Broker # 2	2011-06-15 00:01:18	2011-06-15 00:01:33	P001	192.168. []	1364	NDMCF2	1.1.1.1	52231	4,083
8	Exchange #2	2011-06-15 00:01:19	2011-06-15 00:01:19	P001	192.168. []	21	FTPDN1	2.2.2.2	51491	0
9	Exchange #3	2011-06-15 00:01:19	2011-06-15 00:01:19	P001	192.168. []	21	FTPDN1	3.3.3.3	51492	0

Looks Pretty... but....



- A lot of work involved!
- Needed information is in many places
- Value add: consolidation is good for the organization

Trending and Workload Monitoring

- Who has the most problems?
- Who is using resources?
- Where does the money go?
- Run IT as a business!




Trending Reports

TCP Connections by Remote Host
 Stack ID : P001
 Sorted By : Duplicate Acks Desc
 Showing Entries 1 - 50
 Using Daily Database
 From:2011-03-15 To:2011-03-18



Stack ID Local Host	Remote Host	Number of Connections	Total Bytes In	Total Bytes Out	Total Duplicate Acks	Total Retransmits	Total Good Terminations
Bank # 1	1.2.3.21	1K (74.75%)	21M (40.69%)	635M (78.17%)	4K (69.81%)	46 (41.81%)	1K (74.75%)
Bank # 1	1.2.3.75	341 (16.06%)	5M (10.5%)	68M (8.43%)	899 (13.36%)	22 (20.0%)	341 (16.06%)
Bank # 1	1.2.3.17	20 (0.94%)	1K (0.0%)	2M (0.32%)	481 (7.15%)	3 (2.72%)	20 (0.94%)
Bank # 1	1.2.3.14	80 (3.76%)	854K (1.62%)	102M (12.64%)	381 (5.66%)	32 (29.09%)	80 (3.76%)
Bank # 1	1.2.3.18	44 (2.07%)	3K (0.0%)	2M (0.36%)	210 (3.12%)	0 (0.0%)	44 (2.07%)
Bank # 1	1.2.3.43	51 (2.4%)	24M (47.15%)	488K (0.06%)	59 (0.87%)	7 (6.36%)	51 (2.4%)
All Hosts	-	2K	52M	812M	6K	110	2K

TCP Connections by Remote Host Sorted By : Duplicate Acks

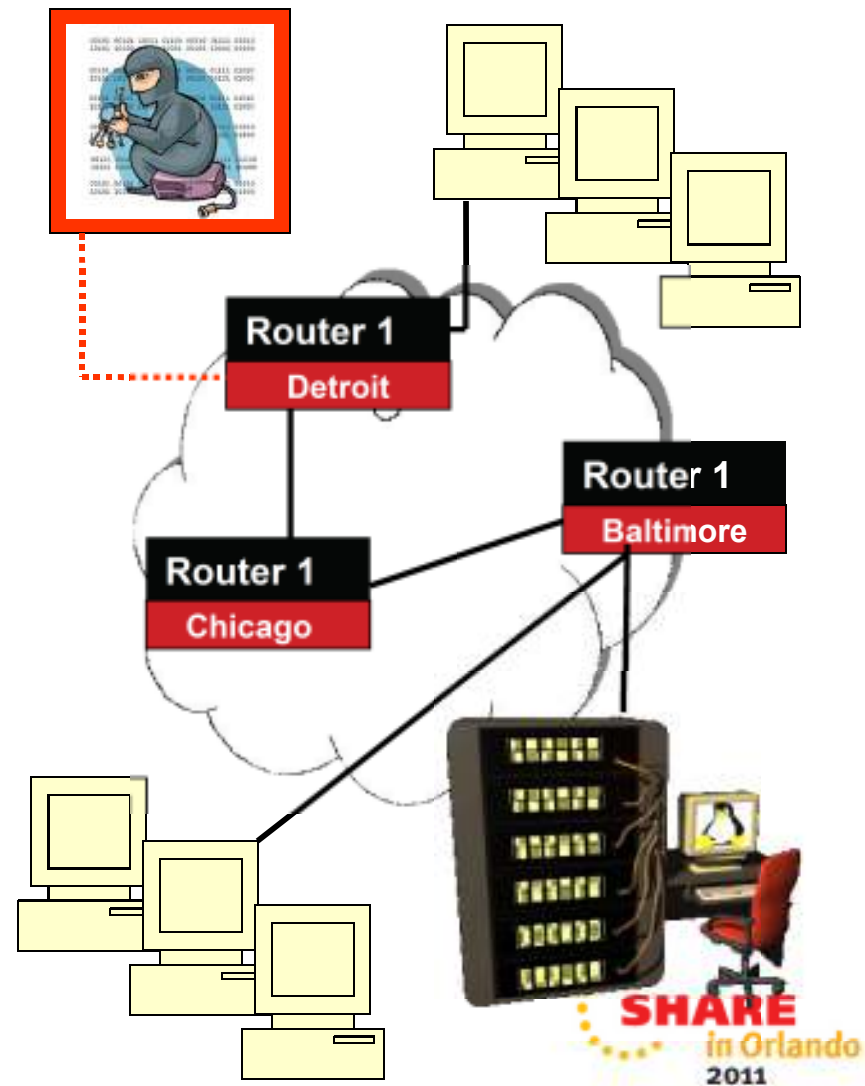


1.2.3.21	1.2.3.75	1.2.3.17	1.2.3.14	1.2.3.18	Other
----------	----------	----------	----------	----------	-------

Keeping the Network Available

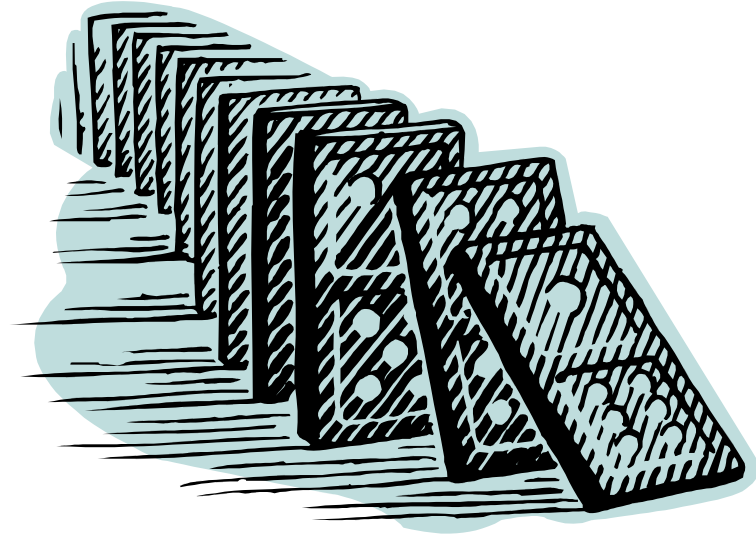
- Mainframes
- Switches
- Routers
- Servers
- Applications

- All have to be monitored in different ways
- Fallback strategies devised



Monitor Unavailability

- Metrics
 - Times unavailable
 - Duration of unavailability
 - Unavailable from where?
- Correlate unavailability with other resources
- Can have domino effect (one resource going down impacts another)



**Service Delivery
Impact !!!**

How we do Availability Checking



Group	Host Name	Source IP	Address Monitored	Available	Last Monitored
	1.2.3.4	127.0.0.1	1.2.3.4	No	2006-06-09 13:21:18.0
231Group	1.2.3.4	1.2.3.4	1.2.3.4	Yes	2006-06-09 13:22:19.0
232Group	1.2.3.4		1.2.3.4	Yes	2006-06-09 13:21:17.0

We use Availability Checker, a real-time monitor checking the availability of routers, switches, servers and any other networked devices that is in communication with the mainframe.

ICMP (Ping) is used to perform this activity. Availability Checker requests are submitted to the mainframe to Ping the remote device. Mainframe to end-device connectivity monitor generates SNMP Traps that are shipped to the NETCOOL Server. Unavailable Resource is considered a critical event.

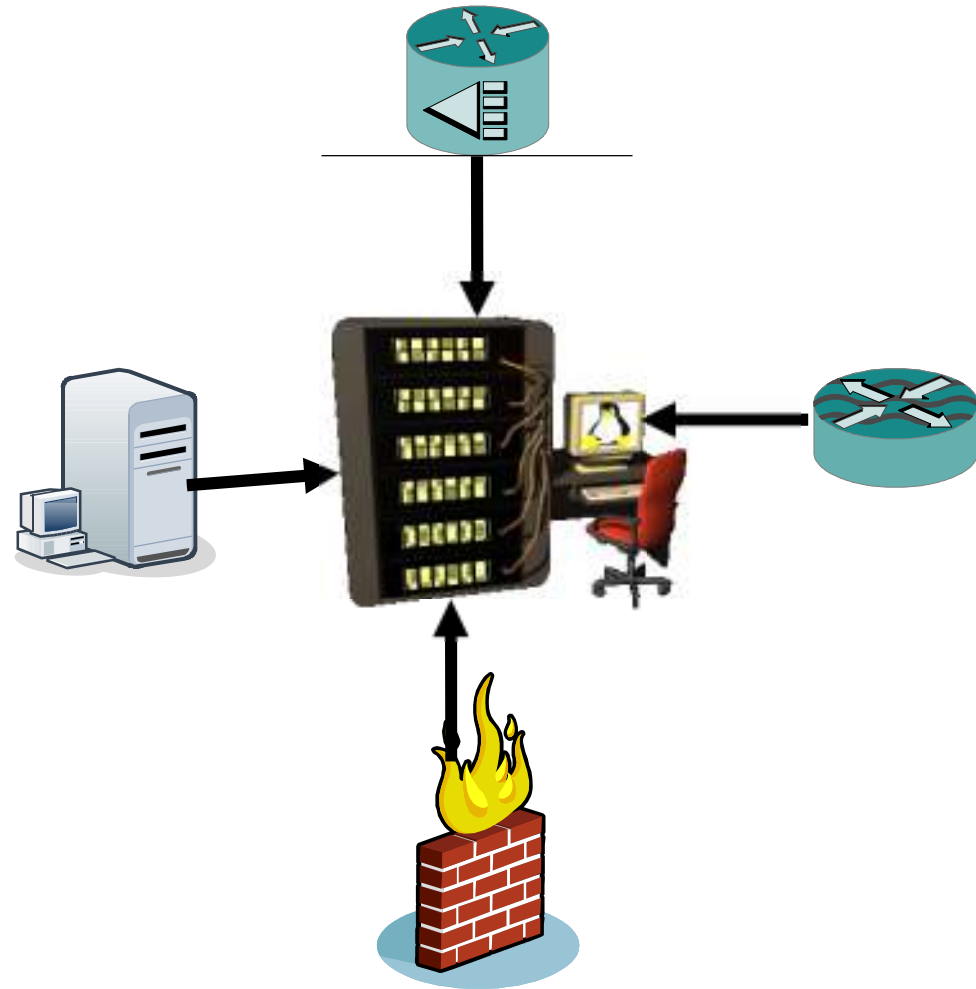
This tool provides access to historical data about network device response time (hourly, daily, weekly, yearly) and unavailable resources.



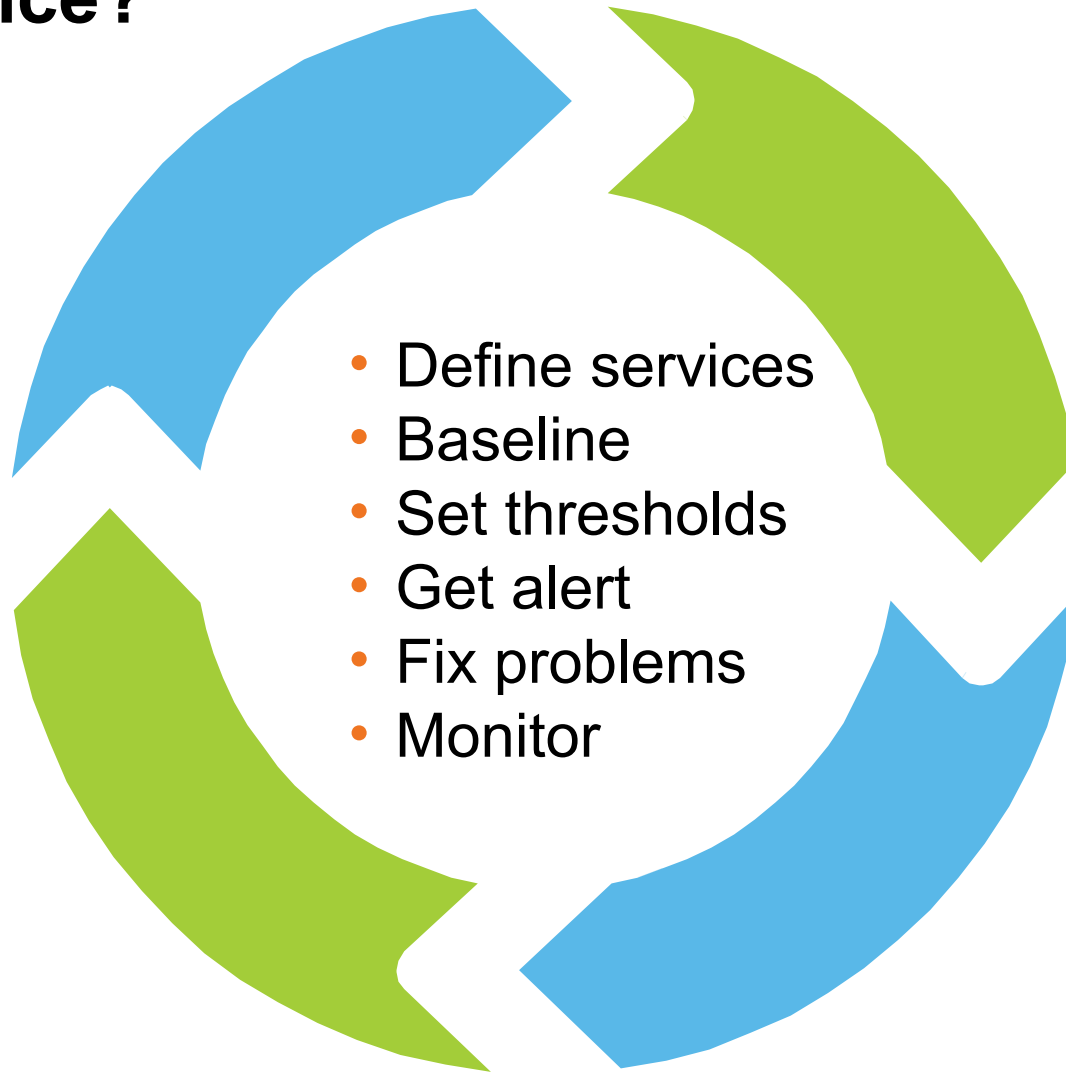
What do we monitor?

- Routers
- Switches
- Firewalls
- Application servers

- All must have connectivity to all mainframes.



How do we know if we are providing good service?



Create Baseline



- Baselining your network is a crucial task
- Problem areas: shift changes, weekends, red letter days
- Re-baselining with changes in topology

Set Thresholds



Customers with different profiles require different thresholds .

How do we set thresholds?

- Separate baseline data on a per client basis.
 - Round-trip time
 - Duplicate acknowledgments
 - Retransmissions
 - Bytes in / out
- Apply statistical measurements. For example:
 - Median
 - Maximum
 - 90th percentile

Sample Duplicate Ack Analysis

	Partner 1	Partner 2	Partner 3	Partner 4	Partner 5
Minimum	0	0	0	0	0
Median	1	1	1	1	1
90th percentile	3	1	24	14	11
95th percentile	3	2	40	34	40
98th percentile	6	8	747	98	319
99th percentile	11	164	902	288	648
Maximum	462	2,918	1,371	33,794	14,225
Suggested Warning Threshold	5	5	50	40	50
Suggested Critical Threshold	25	175	1,000	300	650



Produces Warning Definitions

Sample Client Configuration

- Each client is different
- Want to warn at different levels

IDENTIFIER	Partner 1
IP-ADDRESS4	123.456.*.*
IP-ADDRESS6	NA
LOCAL-PORT	*
REMOTE-PORT	*
MONITOR-INTERVAL	60
CONGESTION-WINDOW	5000
ROUND-TRIP-TIME	250
ROUND-TRIP-VARIANCE	2000
BYTES-OUT	-
RETRANSMITS	2
CONNECTION-TERMINATED	N
DUPLICATE-ACKS	4
HUNG	10
STATUS	SYNSENT
LOCAL-WINDOW-0	1
REMOTE-WINDOW-0	1
OUT-OF-ORDER	0

Early Warning System
Inside Products, Inc.


Alerting Fundamentals

Don't cry wolf but...



Don't be asleep at the wheel


Some Words on Statistical Soundness



- Why these stats?
- 90th, 95th percentiles show outliers

Sample Size

- Too little data is no good
- Garbage in, garbage out



- Example of small sample
- Only have 3 people, with a height of 5 feet, then graph will be very skewed.

How we got baseline data

- Connection Log (NMI)
- Can also use SMF records
- Data collection / manipulation / storage are quite large issues



Three Step Process

- First, figure out what you are providing currently, and deal with obvious problems.
- Second, once you know what is possible/reasonable, negotiate with customers and get their agreement on what service they want/require.
- Third, if higher capacity required, then negotiate with customer.
- Note: We have actual SLAs to measure against, not just internally generated SLOs (Service Level Objectives).

Getting Alerts

- Sample alerts

2010-06-07 01:30:17.12	ITSWA201W	MY BANK	CON WINDOW	LT	5,000
2010-06-07 01:30:57.14	ITSWA205W	MY BANK	RETRANSMITS	GT	10
2010-06-07 08:30:11.25	ITSWA202W	MY BANK	RD TRIP TIME	GT	100

- Alert correlation

- Goal is to send to NetCool
- SNMP traps
- Correlate with other activity

DTCC Closes the Markets

- Times of the day call for extra capacity
- Do we have it?
- How do we measure it?



DTCC Runs IT as a Business

- How?
- Run lean
- Constant tuning
- High availability

Be prepared for the future!

