VTAM and HPR Performance Update
Session 3644
Wednesday, February 24, 1999, 1:30 PM

Nancy Gates
IBM Networking Systems
Center, Gaithersburg, MD

Trademarks

The following are trademarks of the International Business Machines, Inc.

ACF/VTAM*
AFP
AIX*
AIX/6000*
AnyNet
APL2*
APPN*
AS/400*
Application Development
BookManager*
CICS
CICS/ESA*
CUA
DB2*
DFSMS
DFSMS/MVS*
DFSMSdfp
DFSMSdss
DFSMShsm
DFSMSrmm
ESCON*
ES/9000*
GDDM*
geoManager*
Hardware Configuration Definition
Hiperbatch
Hiperspace
IBM*
IMS
IMS/ESA*
Language Environment*
MVS/DFP

The following are trademarks of other corporations:

C++
CMA-CONNECTION
DCE
Endevor
FINGRAPH II/M
HP/UX
Intel
Interlink
Librarian
MAC
Motif
NetWare
NFS
Novell
OSF
Panvalet
POSIX
SAS/GRAPH**
Sun
UNIX**
WINDOWS
XPG4
X/Open

Acronyms

VTAM - Virtual Telecommunications Access Method
HPR - High Performance Routing
NPM - NetView Performance Monitor
NGMF - NetView Graphics Monitor Facility
RTP - Rapid Transport Protocol
ANR - Automatic Network Routing
NCP - Network Control Program
CMIP - Common Management Information Protocol
APPN - Advanced Peer to Peer Networking
COS - Class of Service
SNA - Systems Network Architecture
SNI - SNA Network Interconnection
CDS - Central Directory Server
AHHC - APPN Host to Host Channel
APPNTAM - APPN Topology and Accounting Management
NN - Network Node
EN - End Node
ICN - Interchange Node
MDH - Migration Data Host
FFST - First Failure Support Technology
NETID - Network Identifier
RODM - Resource Object Data Manager
VIT - VTAM Internal Trace
VSE/ESA - Virtual Storage Extended/Enterprise Systems Architecutre
VM/ESA - Virtual Machine/Enterprise Systems Architecture
MVS/ESA - Multiple Virtual Storage/Enterprise Systems Architecture
ATM - Asynchronous Transfer Mode
OS/390 - Operating System for S/390
CS/390 - Communication Server for System/390

Agenda

- *VTAM V4R4 Performance Information*
  - *VTAM V4R4 Benchmark Results*
  - *VTAM V4R4 performance related enhancements*
  - *Storage and Cycle Tuning general guidelines*
- *HPR Performance Information*
  - *Review of HPR performance Information*
  - *Recent HPR benchmark performance information*
  - *HPR tuning general guidelines*

VTAM V4R4

- *VTAM V4R4 has several functions directed at improved availability with HPR and further exploitation of high speed networking*
  - *High Performance Data Transfer (MPC+)*
  - *Native ATM with OSA-2 adapter*
  - *Multi-node Persistent Sessions*
  - *HPR available in many configurations*
  - *Miscellaneous enhancements*
- *Extensive VTAM V4R4 benchmarks were conducted in Research Triangle Park, NC in early 1997*
  - *Results documented in paper available via:*
    - *System Center Flash www-1.ibm.com/support/techdocs/atsflash.nsf*
    - *Support Page www.software.ibm.com/network/commserver*

VTAM V4R4 High Performance Data Transfer

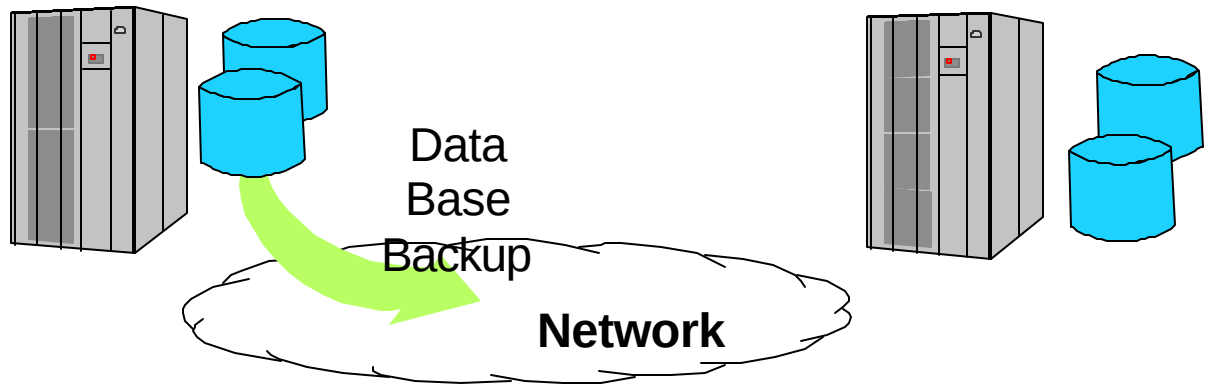

- **MPC+ is a channel protocol first introduced with VTAM V4R4 (OS/390 V1R3 and higher)**
- **Full duplex protocol, exploiting Communication Storage Manager (CSM) and Seldom Ending Channel Programs**
- **Used when connections traverse VTAM V4R4 high bandwidth HPR network attachments**
  - **S/390 OSA connected to native ATM network**
  - **VTAM Host to Host Channel Connections**
  - **XCF links between processors in a Sysplex**
  - **VTAM to 2216 or 3746 MAE MPC+ connections**
- **Can increase channel utilization, decrease CPU utilization and improve throughput**

Data Moves Prior to MPC+

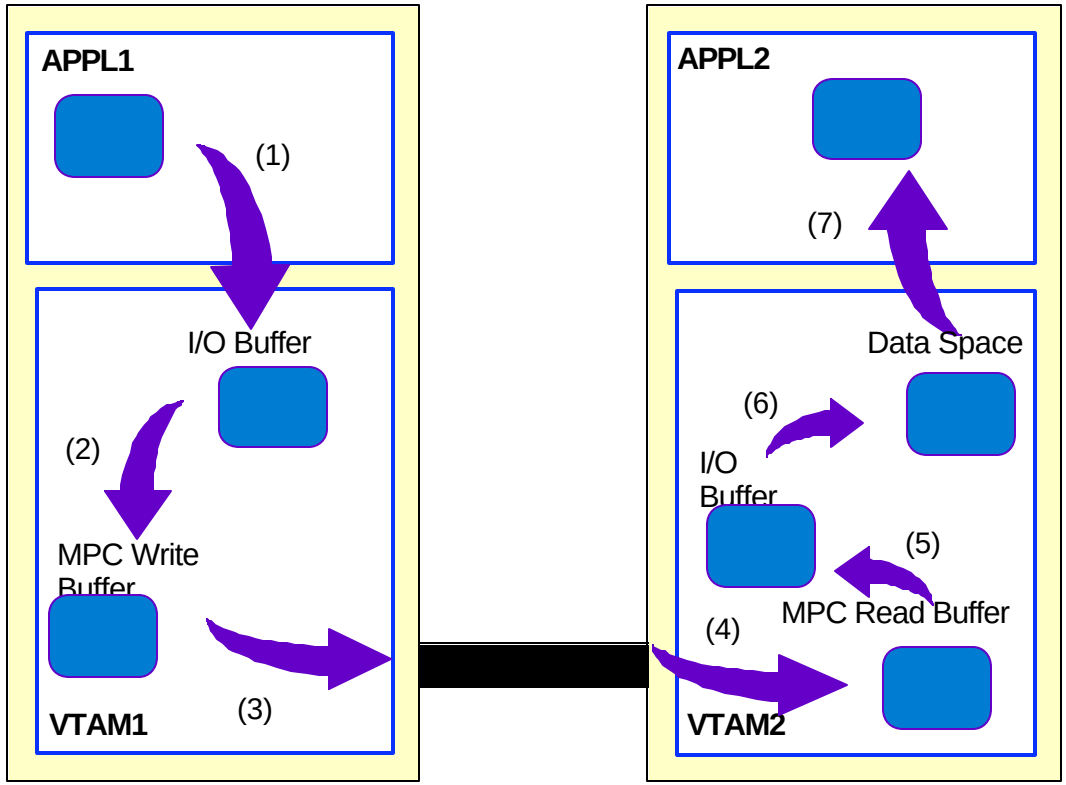

(1) Move data from APPL's buffers to IO buffers

(2) Move data from IO buffers to MPC write buffers

(3) Write data over channel

(4) Read data into MPC read buffers

(5) Move data from MPC read buffers to IO buffers

(6) If RECEIVE not outstanding then move data from IO buffers into Data Space storage

(7) Move data from Data Space storage to APPL's buffers once RECEIVE issued

Data Moves with MPC+

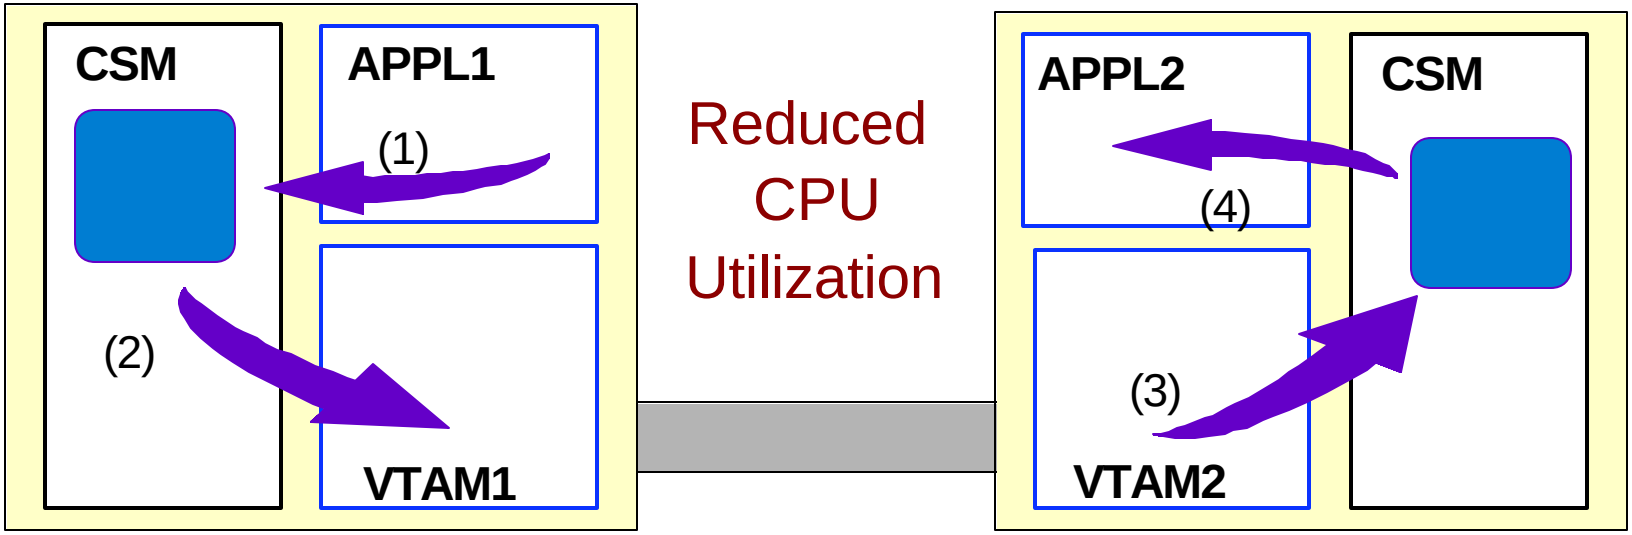

(1) Move data from APPL's buffers into CSM buffers

(2) MPC+ accesses CSM buffers to write data over channel

(3) MPC+ reads data into CSM data space buffers waiting on RECEIVE

(4) Move data from CSM Data Space buffers to APPL's buffers once RECEIVE issued

- **CSM code ships as part of VTAM V4R4**
  - **Runs as a separate MVS task**
  - **CSM storage in Data Space or ECSA**

Data Moves with MPC+ and HPDT

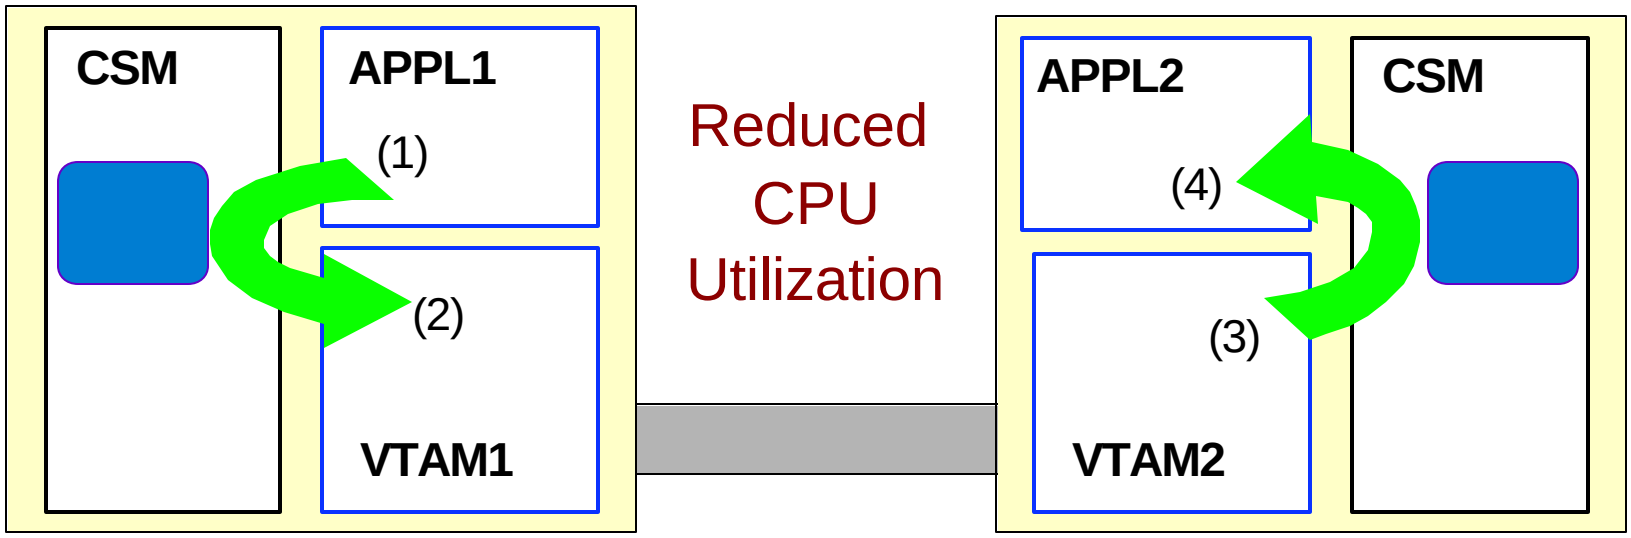

(1) APPL writes data into CSM buffers and passes ownership of buffers to VTAM

(2) MPC+ accesses CSM buffers to write over channel

(3) MPC+ reads data into CSM data space buffers waiting on RECEIVE

(4) VTAM passes ownership of CSM buffers to APPL once RECEIVE is issued

- **VTAM/APPC applications can use High Performance Data Transfer (HPDT) application program interface**
  - **Additional performance improvements for large blocks because of less data moves**

VTAM V4R4 HPDT and MPC+ Benchmarks

The following test cases were run using the NetMarks program:

- ***44 HPR***
  - ***Netmarks with HPDT API -- CSM storage in ECSA***
  - ***VTAM V4R4 to VTAM V4R4 MPC+ ESCON connections***
- ***44 HPR Non-HPDT API***
  - ***Netmarks using VTAM/APPC - CSM storage in ECSA***
  - ***VTAM V4R4 to VTAM V4R4 MPC+ ESCON connections***
- ***44 HPR Non-HPDT API DS***
  - ***Netmarks using VTAM/APPC - CSM storage in DataSpace***
  - ***VTAM V4R4 to VTAM V4R4 MPC+ ESCON connections***
- ***44 HPR (No MPC+)***
  - ***Netmarks using VTAM/APPC***
  - ***VTAM V4R4 to VTAM V4R4 AHHC ESCON connections (not MPC+)***
- ***43 HPR***
  - ***Netmarks using VTAM/APPC***
  - ***VTAM V4R3 to VTAM V4R3 AHHC ESCON connections***
- ***44 Subarea MPC***
  - ***Netmarks using VTAM/APPC***
  - ***VTAM V4R4 to VTAM V4R4 MPC (Subarea) ESCON connections***
- ***43 Subarea MPC***
  - ***Netmarks using VTAM/APPC***
  - ***VTAM V4R3 to VTAM V4R3 MPC (Subarea) ESCON connections***

VTAM V4R4 First Measurement Set

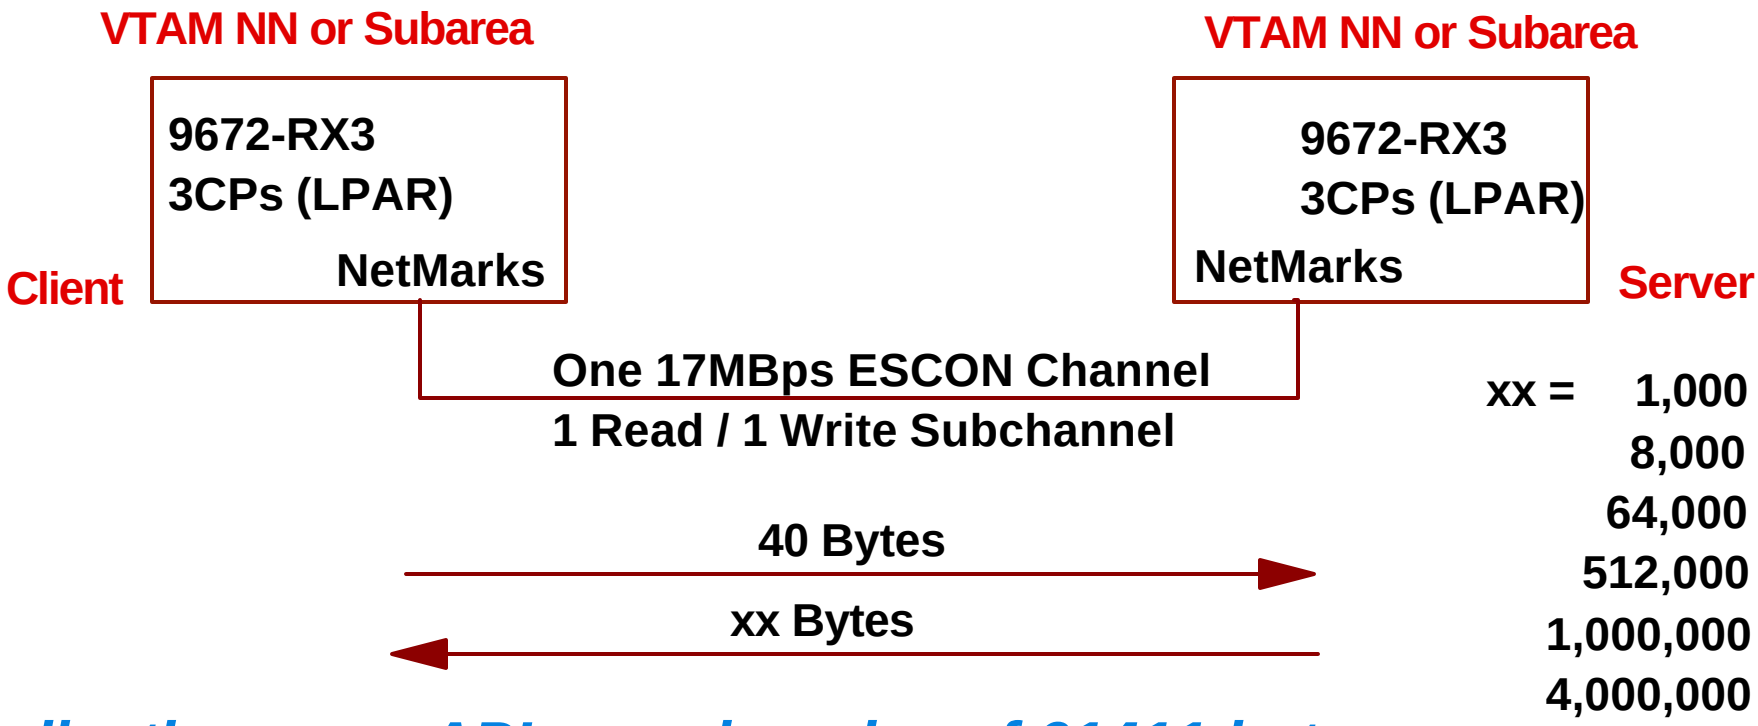

- *Application uses API crossing size of 61411 bytes*
- *RUsize in BIND is 60K*
- *VTAM IOBUF set to largest size (3992 bytes)*
- *Pacing windows are set to 63 (adaptive session pacing also applies)*
- *DELAY = 0 (no coattailing delay)*
- *CSM initial and expansion sizes set to 500 and 200*

MPC+ With/Without HPDT Measurements

File Size	Parameters	44 HPR	44 HPR non-HPDT API	43 HPR
1,000	Thruput Bytes/Sec	205.637	224,416	226,540
	CPU % Client	18.87%	19.00%	22.28%
	Server	19.77%	19.66%	24.37%
	CPU Sec/Byte Client	2.752E-6	2.540E-6	2.950E-6
	Server	2.884E-6	2.628E-6	3.227E-6
8,000	Thruput Bytes/Sec	1,472,436	1,556,596	1,527,417
	CPU % Client	17.90%	17.86%	21.00%
	Server	18.33%	18.21%	22.76%
	CPU Sec/Byte Client	3.647E-7	3.442E-7	4.125E-7
	Server	3.735E-7	3.510E-7	4.470E-7
64,000	Thruput Bytes/Sec	6,690,790	6,164,534	5,514,138
	CPU % Client	16.89%	17.00%	19.76%
	Server	15.10%	15.59%	20.11%
	CPU Sec/Byte Client	7.573E-8	8.273E-8	1.075E-7
	Server	6.771E-8	7.587E-8	1.094E-7
512,000	Thruput Bytes/Sec	13,540,844	13,245,419	10,777,640
	CPU % Client	15.99%	19.74%	21.82%
	Server	12.23%	16.54%	24.52%
	CPU Sec/Byte Client	3.543E-8	4.471E-8	6.074E-8
	Server	2.710E-8	3.746E-8	6.825E-8
1,000,000	Thruput Bytes/Sec	14,223,854	14,041,280	11,502,200
	CPU % Client	15.71%	20.01%	22.23%
	Server	11.98%	16.68%	24.92%
	CPU Sec/Byte Client	3.313E-8	4.275E-8	5.798E-8
	Server	2.527E-8	3.564E-8	6.500E-8
4,000,000	Thruput Bytes/Sec	15,053,724	14,956,889	12,134,400
	CPU % Client	15.93%	20.25%	22.89%
	Server	11.83%	19.54%	25.52%
	CPU Sec/Byte Client	3.175E-8	4.062E-8	5.659E-8
	Server	2.358E-8	3.919E-8	6.309E-8

VTAM V4R4/V4R3 AHHC Measurements

File Size	Parameters	44 HPR (no MPC+)	43 HPR
1,000	Thruput bytes/sec	243,171	226,540
	CPU % Client	19.80%	22.28%
	Server	19.49%	24.37%
	CPU Sec/Byte Client	2.443E-6	2.950E-6
	Server	2.404E-6	3.227E-6
8,000	Thruput bytes/sec	1,591,739	1,527,417
	CPU % Client	18.31%	21.00%
	Server	18.34%	22.76%
	CPU Sec/Byte Client	3.451E-7	4.125E-7
	Server	3.457E-7	4.470E-7
64,000	Thruput bytes/sec	5,522,700	5.514,138
	CPU % Client	17.22%	19.76%
	Server	17.03%	20.11%
	CPU Sec/Byte Client	9.354E-8	1.075E-7
	Server	9.251E-8	1.094E-7
512,000	Thruput bytes/sec	10,210,500	10,777,640
	CPU % Client	18.74%	21.82%
	Server	20.83%	24.52%
	CPU Sec/Byte Client	5.506E-8	6.074E-8
	Server	6.120E-8	6.825E-8
1,000,000	Thruput bytes/sec	10,874,213	11,502,200
	CPU % Client	19.16%	22.23%
	Server	21.73%	24.92%
	CPU Sec/Byte Client	5.286E-8	5.798E-8
	Server	5.995E-8	6.500E-8
4,000,000	Thruput bytes/sec	11,815,633	12,134,400
	CPU % Client	20.15%	22.89%
	Server	23.05%	25.52%
	CPU Sec/Byte Client	5.116E-8	5.659E-8
	Server	5.852E-8	6.309E-8

VTAM V4R4/V4R3 Subarea MPC Measurements

File Size	Parameters	44 Subarea MPC	43 Subarea MPC
1,000	Thruput bytes/sec	314,542	320,555
	CPU % Client	18.30%	17.97%
	Server	19.11%	18.75%
	CPU Sec/Byte Client	1.745E-6	1.682E-6
	Server	1.823E-6	1.755E-6
8,000	Thruput bytes/sec	2,008,189	2,047,008
	CPU % Client	16.77%	16.50%
	Server	17.46%	17.36%
	CPU Sec/Byte Client	2.505E-7	2.418E-7
	Server	2.608E-7	2.544E-7
64,000	Thruput bytes/sec	6.171,537	6,239,905
	CPU % Client	15.09%	14.90%
	Server	15.19%	15.09%
	CPU Sec/Byte Client	7.335E-8	7.164E-8
	Server	7.384E-8	7.255E-8
512,000	Thruput bytes/sec	10,889,385	10,894,619
	CPU % Client	17.01%	16.60%
	Server	17.10%	16.89%
	CPU Sec/Byte Client	4.686E-8	4.571E-8
	Server	4.711E-8	4.651E-8
1,000,000	Thruput bytes/sec	11,582,450	11,560,279
	CPU % Client	17.62%	17.25%
	Server	17.59%	17.40%
	CPU Sec/Byte Client	4.564E-8	4.477E-8
	Server	4.556E-8	4.515E-8
4,000,000	Thruput bytes/sec	12,296,736	12,258,663
	CPU % Client	18.59%	18.10%
	Server	18.05%	17.79%
	CPU Sec/Byte Client	4.535E-8	4.430E-8
	Server	4.404E-8	4.354E-8

Conclusions - Measurement Set I

- **Conclusions about VTAM V4R4 MPC+**
  - *With larger blocksizes, HPDT Application Program Interface (API) significantly increases throughput*
  - *With larger blocksizes, HPDT API significantly reduces VTAM CPU requirements*
  - *With smaller blocksizes, non-HPDT API is slightly better than HPDT API in terms of throughput and VTAM CPU requirements*
  - *VTAM V4R4 MPC+ with or without HPDT API significantly reduces CPU requirements as compared to VTAM V4R3*
  - *VTAM V4R4 MPC+ with or without HPDT API significantly increases throughput with larger blocksizes as compared to VTAM V4R3*
    - ***With smaller blocksizes, VTAM V4R3 HPR throughput is minimally better than VTAM V4R4***

Conclusions - Measurement Set I

- Comparisons of VTAM V4R4 HPR (not MPC+) with VTAM V4R3 HPR:
  - Throughput is about the same
  - VTAM V4R4 reduces CPU requirements
- Comparisons of VTAM V4R4 Subarea MPC with VTAM V4R3 Subarea MPC:
  - VTAM V4R4 performance is slightly degraded in terms of throughput and CPU requirements due to additional functions added with VTAM V4R4
- VTAM V4R4 MPC+ significantly increases throughput and reduces CPU requirements for larger blocksizes as compared to Subarea MPC
- Maximum channel utilization achieved with VTAM V4R4 MPC+ was 95% as compared to maximum of 78-80% with VTAM V4R3

VTAM V4R4 Second Measurement Set

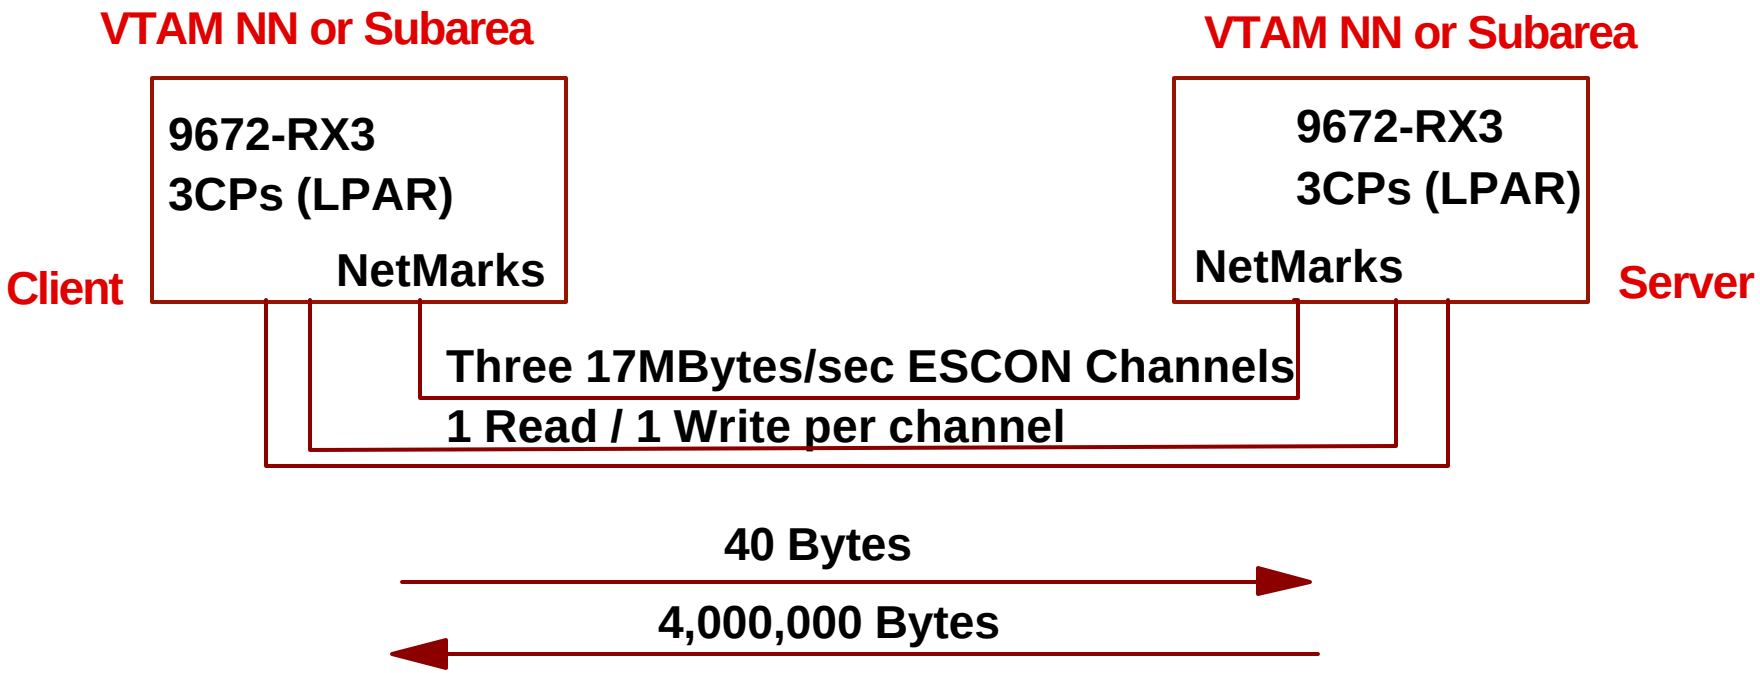

- ***Connections between hosts are:***
  - ***HPDT API and MPC+ (Hosts are VTAM V4R4)***
  - ***Non-HPDT API and MPC+ - either CSM DataSpace or ECSA (Hosts are VTAM V4R4)***
  - ***HPR (Hosts are VTAM V4R3)***
  - ***Subarea MPC (Hosts are both V4R4 or both V4R3)***

Second Measurement Set Results

	Thruput	CPU Utilization		CPU Seconds/Byte	
	Bytes/Sec	Client	Server	Client	Server
44 HPR	41,513,637	43.09%	31.85%	3.114E-8	2.302E-8
44 HPR non-HPDT API	39,056,113	90.76%	50.13%	6.972E-8	3.851E-8
44 HPR non-HPDT API Dataspace	40,379,001	61.65%	49.05%	4.580E-8	3.644E-8
44 HPR (no MPC+)	36,095,776	83.73%	74.52%	6.959E-8	6.194E-8
43 HPR	35.824,934	85.02%	74.61%	7.120E-8	6.248E-8
44 Subarea MPC	37,123,139	75.79%	55.96%	6.125E-8	4.522E-8
43 Subarea MPC	37,038,682	74.54%	54.50%	6.038E-8	4.414E-8

Second Measurement Set Conclusions

- *VTAM V4R4 as compared to VTAM V4R3 HPR:*

	Thruput % Improvement	CPU Utilization (Normalized) % Improvement	
		Client	Server
44 MPC+ HPDT API	15.9%	56.3%	63.2%
44 MPC+ non-HPDT API	9.0%	2.1%	38.4%
44 MPC+ non-HPDT API CSM Dataspace	12.7%	35.7%	41.7%

- *VTAM V4R4 with MPC+ with or without HPDT API significantly improves throughput while reducing CPU requirements*
- *Recommend using CSM dataspace (the default) rather than ECSA*
- *There is minimal performance difference between VTAM V4R4 with no MPC+ and VTAM V4R3*
- *There is minimal performance difference between VTAM V4R4 and VTAM V4R3 Subarea MPC*

VTAM V4R4 Third Measurement Set

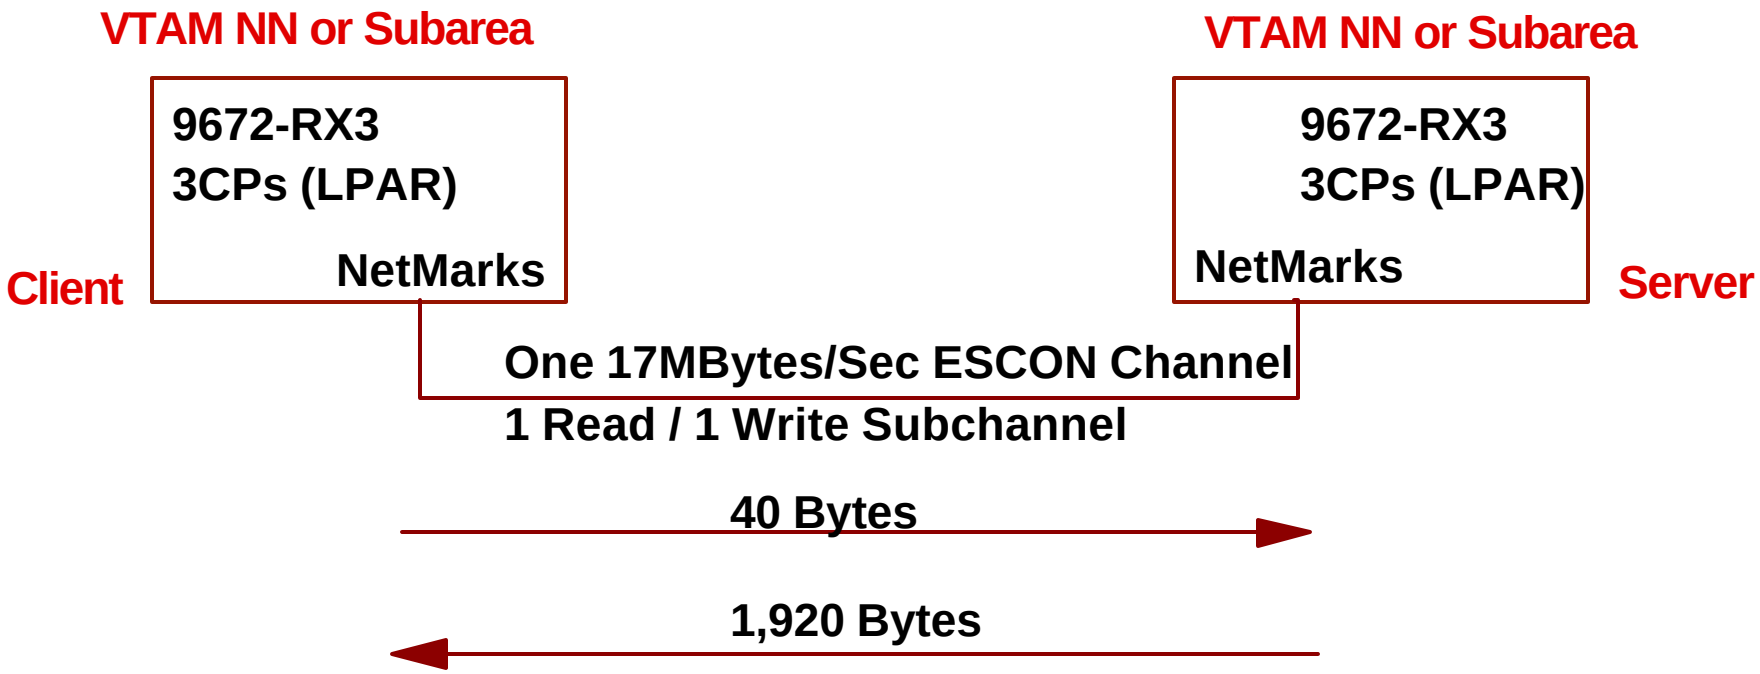

- *Connections between hosts are:*
  - *HPDT API and MPC+ (Hosts are VTAM V4R4)*
  - *Non-HPDT API and MPC+ (Hosts are VTAM V4R4)*
  - *HPR (Hosts are VTAM V4R3)*
  - *Subarea MPC (Hosts are both V4R4 or both V4R3)*

Third Measurement Set Results

	Transactions	CPU Utilization		CPU Seconds/Trans	
	Trans/Sec	Client	Server	Client	Server
44 HPR	925.17	84.77%	83.94%	2.749E-3	2.722E-3
44 HPR non-HPDT API	1017.82	87.41%	84.57%	2.576E-3	2.493E-3
44 HPR (no MPC+)	1076.73	87.50%	87.19%	2.438E-3	2.429E-3
43 HPR	991.08	84.03%	83.40%	2.544E-3	2.525E-3
44 Subarea MPC	1338.23	81.34%	80.67%	1.823EE-3	1.808E-3
43 Subarea MPC	1366.90	79.72%	78.15%	1.750E-3	1.715E-3

Third Measurement Set Conclusions

- *Interactive (small blocksize) applications using the HPDT API may see some degradation (10%) in throughput and some increase (8-9%) in CPU utilization compared to VTAM V4R4 MPC+ without the HPDT API*
- *Interactive applications using MPC+ may see minimal increase (3%) in throughput with the same CPU requirements as compared to VTAM V4R3 HPR*
- *Interactive applications not using MPC+ or HPDT API may see somewhat better throughput (8-9%) while reducing CPU seconds/transaction by 4-5%*
- *VTAM V4R4 Subarea MPC users may see a minimal performance degradation in throughput and CPU utilization as compared to VTAM V4R3 Subarea MPC*

1997 2216 Benchmark MPC+ Observations

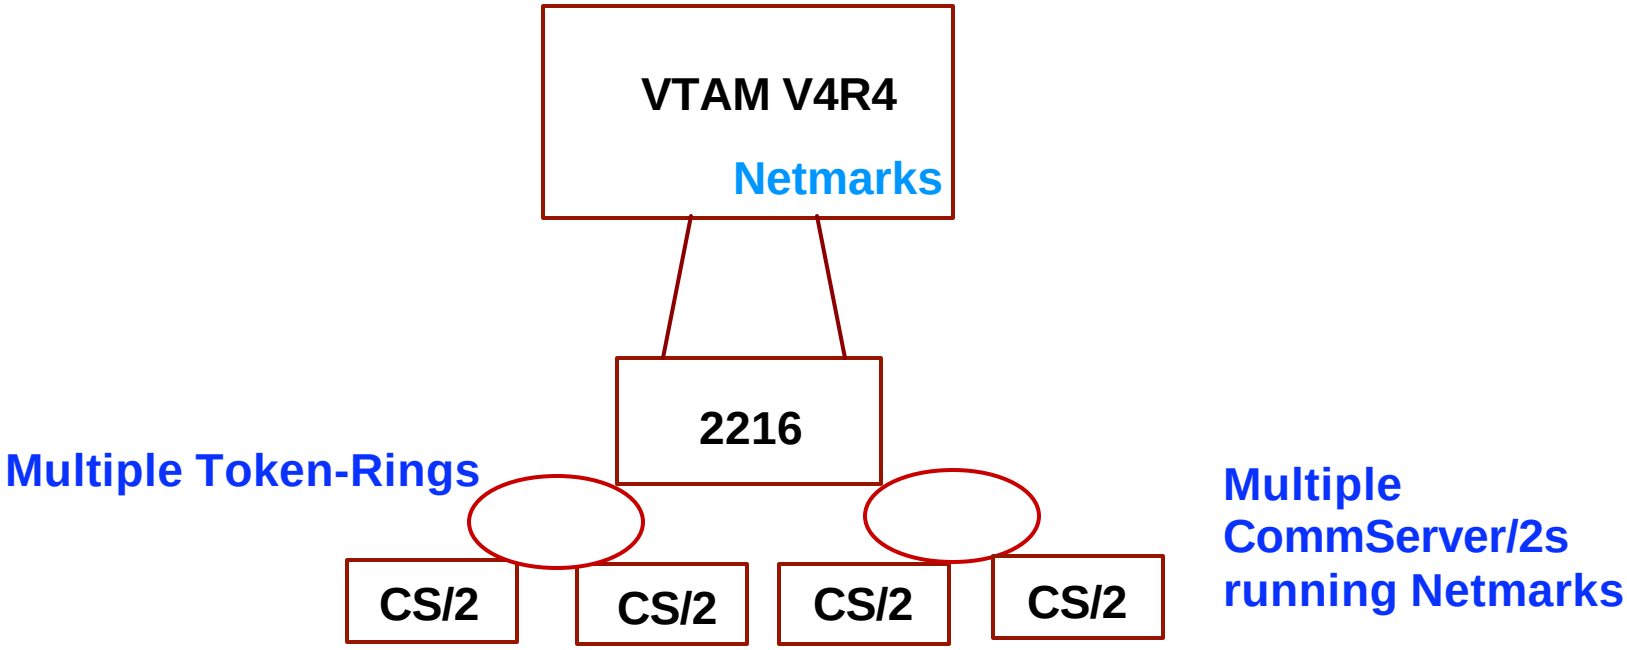

- *Tests run with 16K and 4K RUsizes*
- *Observations roughly similar to VTAM-VTAM MPC+ benchmarks:*
  - *Tests with 16K RUs achieved about twice the throughput as with 4K RUs*
  - *CPU time per byte was more than twice as much with the 4K test than with the 16K test*

VTAM V4R4 ATM Benchmark Configuration

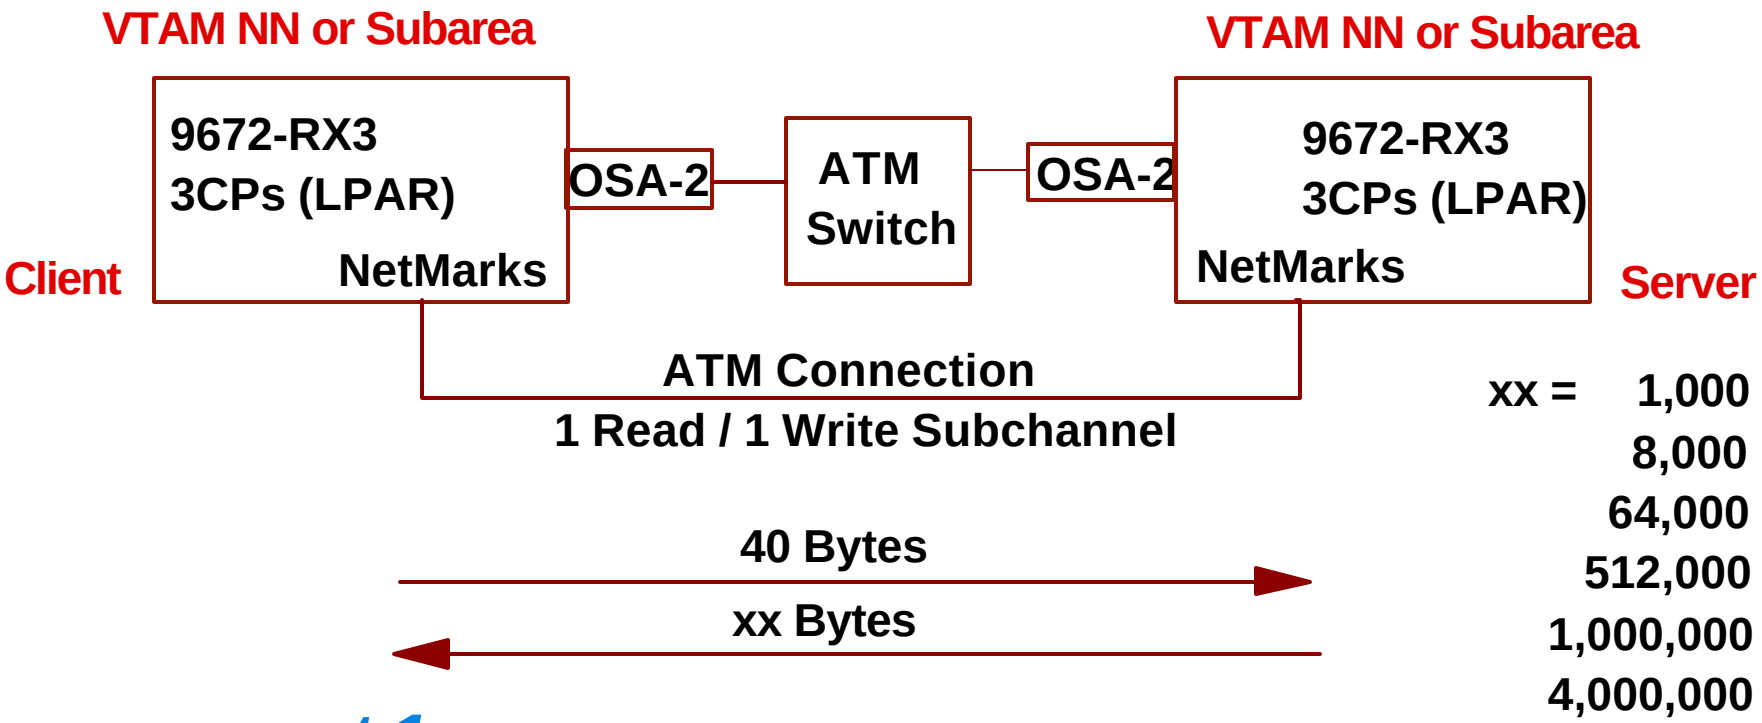

- *Measurement 1*
  - *Application uses API crossing size of 61411 bytes or single API crossing for file sizes less than 61,411 bytes*
  - *RUsize in BIND is 8192 bytes*
- *Measurement 2*
  - *4,000,000 byte file size - RUsize varied between 4 and 16K*
- *CSM storage is in ECSA*

ATM Measurement 1 Results

File Size	Thruput	CPU Utilization		CPU Seconds/Byte	
Bytes	Bytes/Sec	Client	Server	Client	Server
1,000	95,873	10.67%	10.94%	3.339E-6	3.423E-6
8,000	586,820	9.10%	9.11%	4.652E-7	4.657E-7
64,000	1,833,970	10.10%	8.40%	1.652E-7	1.374E-7
512,000	3,973,279	13.70%	10.82%	1.034E-7	8.170E-8
1,000,000	4,490,119	14.82%	11.89%	9.902E-8	7.944E-8
4,000,000	5,208,295	16.32%	14.16%	9.400E-8	8.156E-8

- **Conclusions:**
  - Throughput increases as a function of file size due to latency for acknowledgement contributing a smaller amount to overall time
  - CPU seconds/byte decreases initally but levels off for larger file sizes

ATM Measurement 2 Results

Ruslze	Thruput	CPU Utilization		CPU Seconds/Byte	
Bytes	Bytes/Sec	Client	Server	Client	Server
4,096	4,272,404	21.07%	17.20%	1.479E-7	1.208E-7
7,680	4,990,243	15.33%	12.59%	9.216E-8	7.569E-8
8,192	5,100,781	16.31%	12.60%	9.593E-8	7.411E-8
16,384	5,213,517	13.67%	11.60%	7.866E-8	6.675E-8

- **Conclusions:**
  - *As RUsize increases:*
    - *Throughput increases*
    - *CPU seconds per byte decreases*

VTAM V4R4 Multinode Persistent Sessions

- *Sessions survive failures in VTAM, MVS, or hardware*
- *HPR allows non-disruptive session continuation with application restarted on another processor*
- *VTAM keeps persistent session information in Coupling Facility*
- *Benchmark measures ICN and EN CPU time, elapsed time, and storage*
  - *5,000 and 10,000 sessions*
  - *Session establishment, failure and recovery*

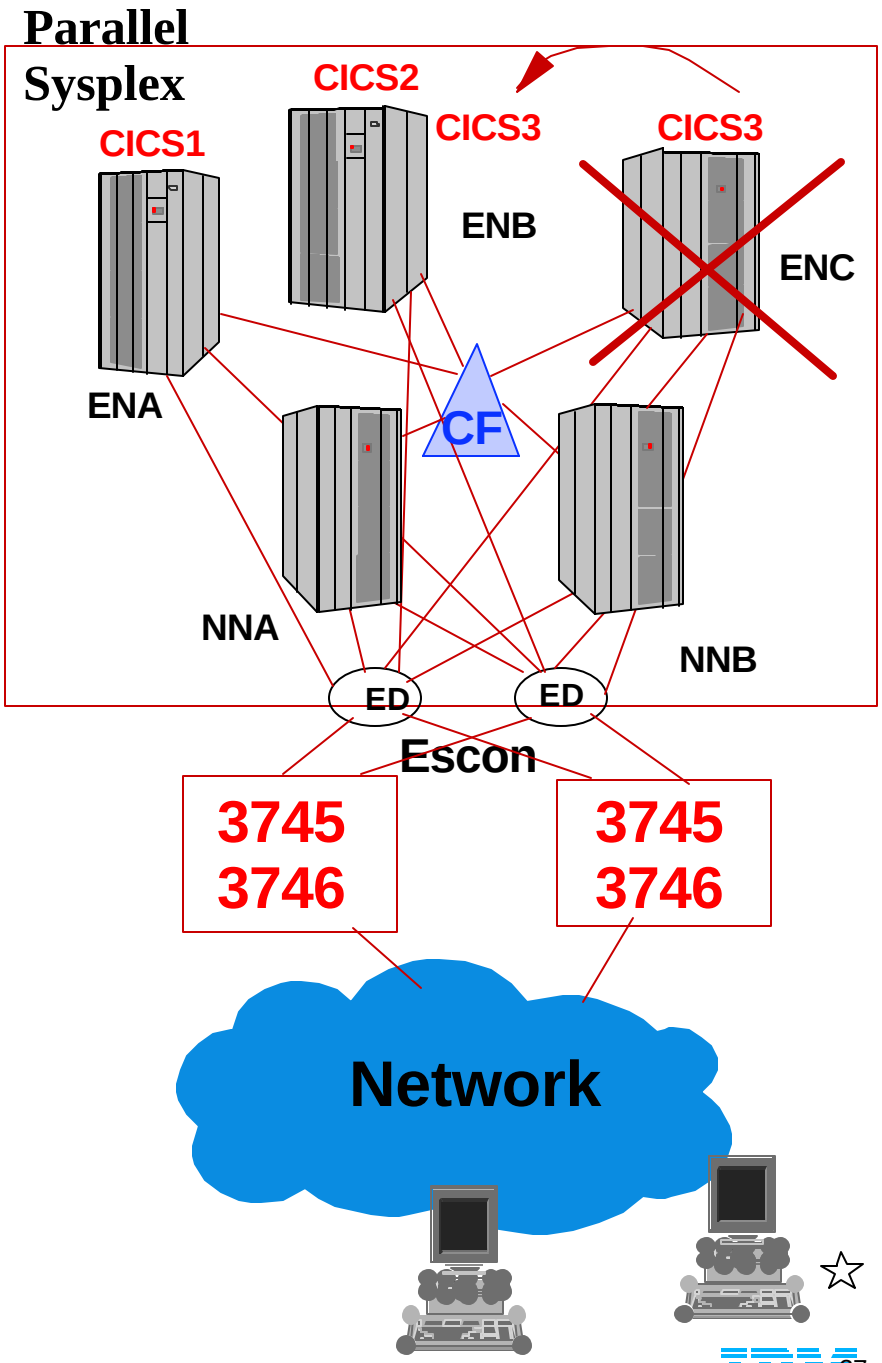

MNPS Session Establishment Results

MNPS versus non-MNPS Session Establishment

	CPU Seconds		Elapsed Time (mm:ss)
	ICN	EN	
5,000 LUs MNPS	105.6	119.1	1:58
5,000 LUs non-MNPS	105.2	93.1	1:57
10,000 LUs MNPS	227.6	254.6	3:24
10,000 LUs non-MNPS	221.3	193.4	3:23

MNPS Storage Requirements

	MNPS Session	SNPS Session	Non-PS Session
ECSA	2.8K	1.3K	1.0K
Private	1.4K	1.1K	1.1K

MNPS Failure and Recovery Results

MNPS Failure and Recovery measurements:

1. 5,000 sessions failed and recovered - recovery host has no other activity
2. 10,000 sessions failed and recovered - recovery host has no other activity
3. 5,000 sessions with send/receive traffic failed and recovered - recovery host has no other activity
4. 10,000 sessions with send/receive traffic failed and recovered - recovery host has no other activity
5. 10,000 sessions failed and recovered - recovery host is 40%-50% busy with other traffic

Failure and Recovery for MNPS Sessions

	Elapsed Time (mm:ss)	VTAM CPU Seconds
5,000 sessions	00:39	41.02
10,000 sessions	01:28	83.97
5,000 sessions w/traffic	01:00	54.00
10,000 sessions w/traffic	03:00	109.93
10,000 sessions w/CPU load	01:24	84.00

VTAM V4R4 APPN Network Measurements

APPN Configuration - Measurements taken for activation, logon, failure, and recovery

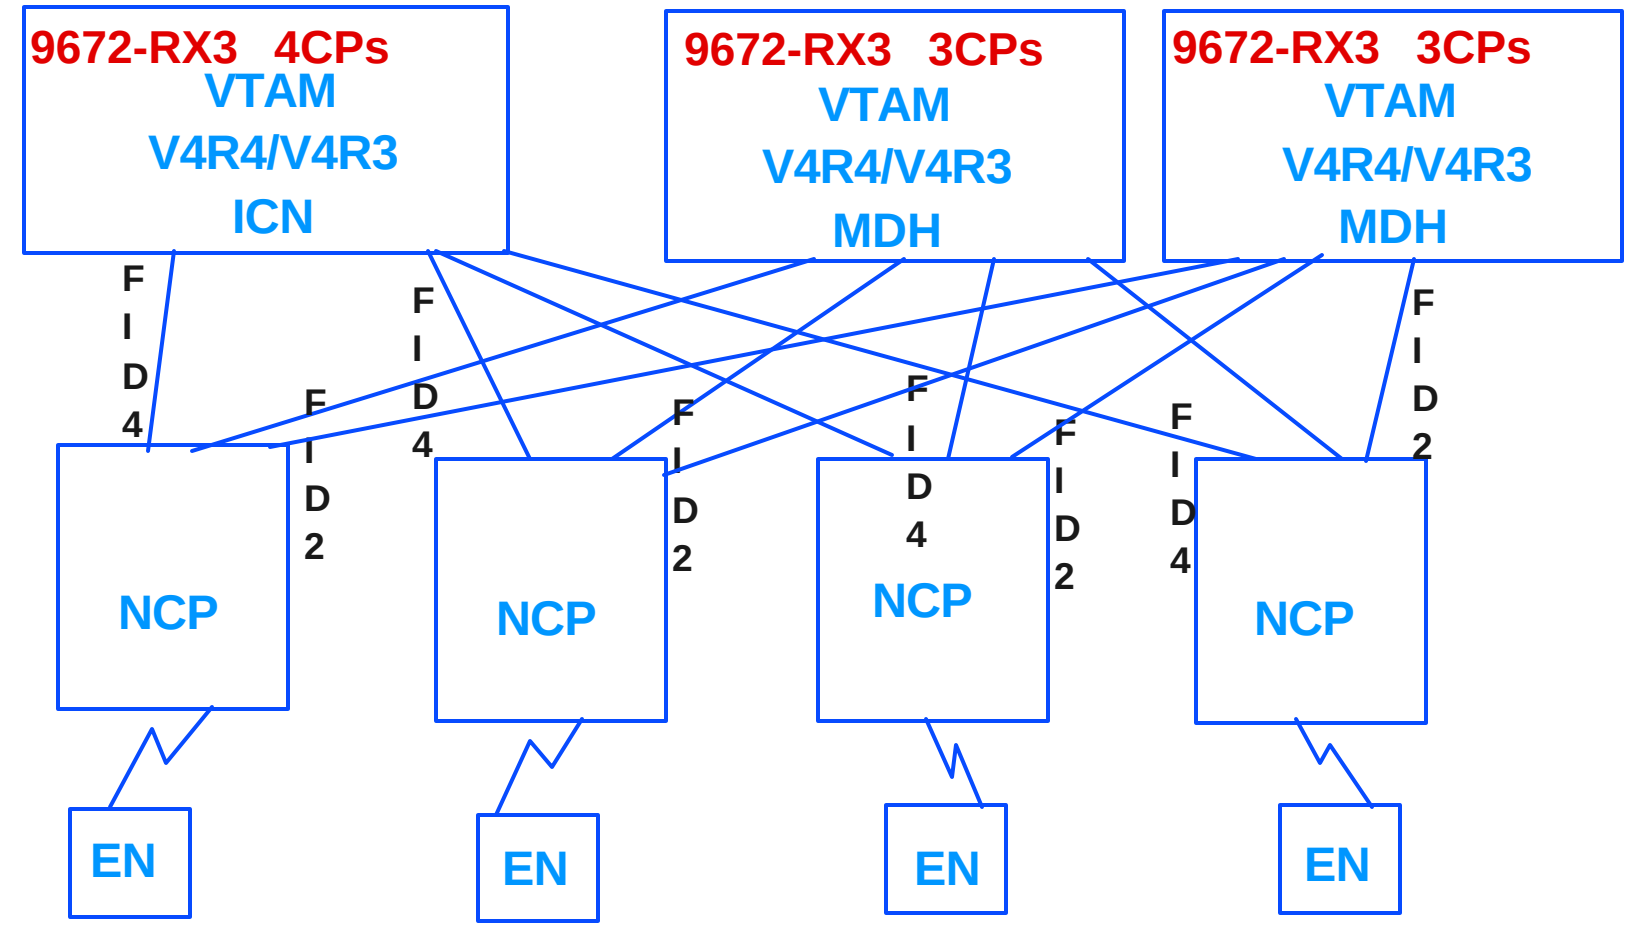

10 NCPs
Each NCP has:
200 ENs
5,000 LUs simulated by TPNS

VTAM 50K APPN Release Comparison

	VTAM V4R4			VTAM V4R3		
	Elapsed Time (mm:ss)	CPU seconds (ICN)	CPU seconds (MDH)	Elapsed Time (mm:ss)	CPU seconds (ICN)	CPU seconds (MDH)
Activation	05:42	283.8	*	06:30	320.2	*
Logon	12.38	970.7	372.2	13:40	1104.1	438.6
Failure/React	04:05	60.9	24.5	04:25	76.8	31.8
Re-logon	03:53	92.0	73.1	04:40	102.6	89.5

Conclusion:

- VTAM V4R4 reduces elapsed and CPU time in all processes as compared to VTAM V4R3

VTAM V4R4 Subarea Network Measurements

Subarea Configuration - Measurements taken for activation, logon, failure, and recovery

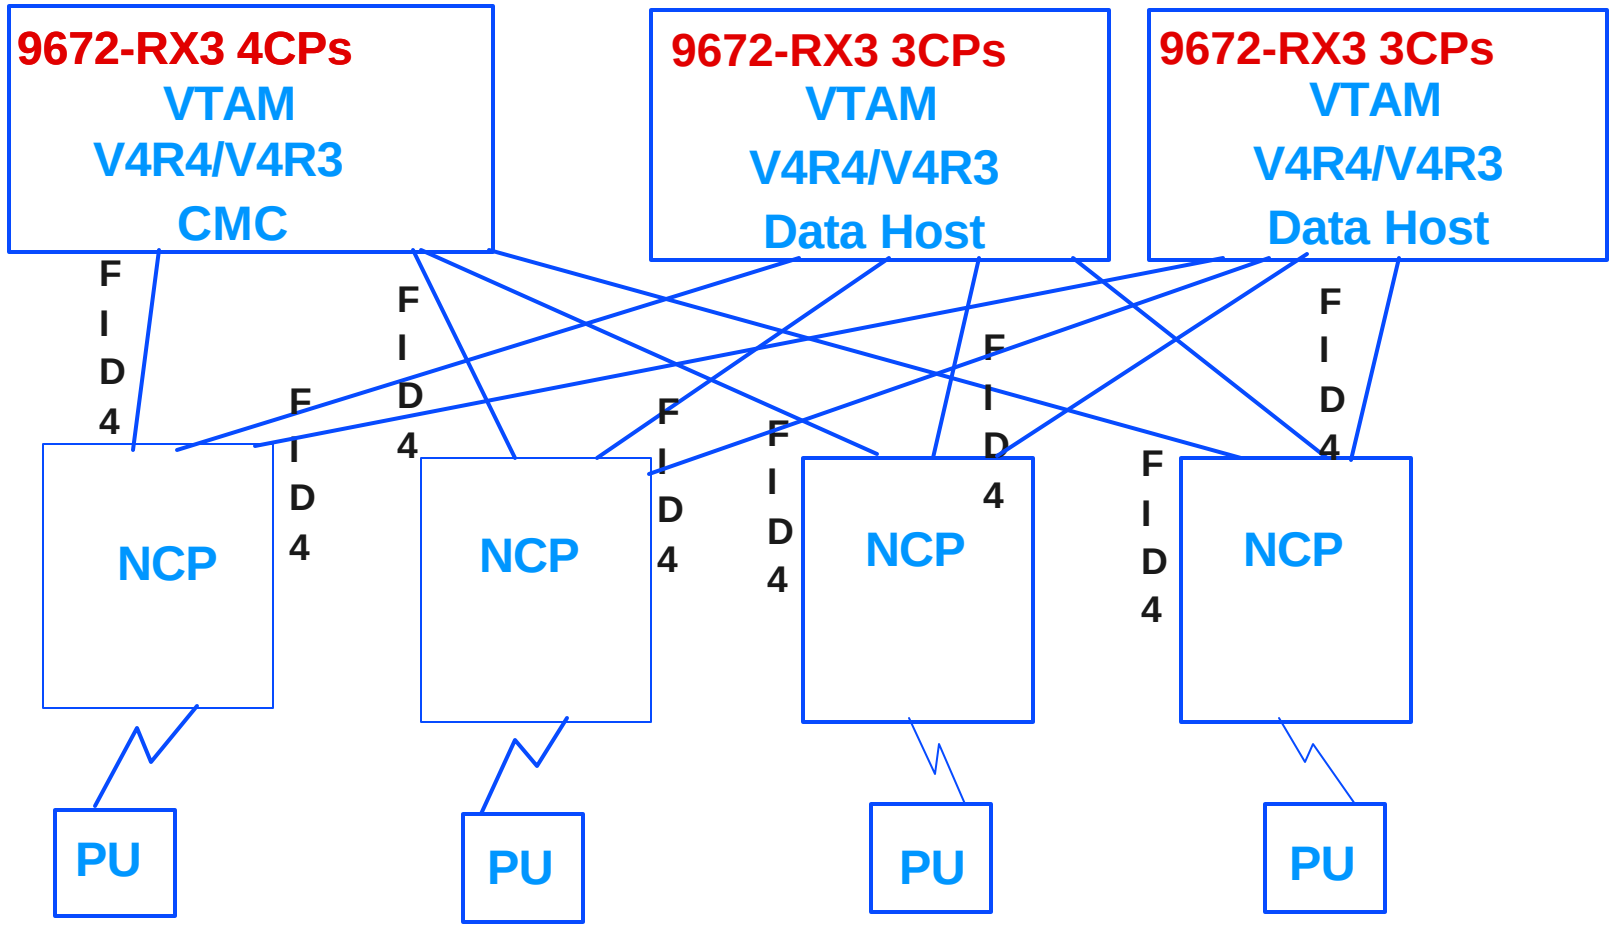

10 NCPs
Each NCP has:
200 PUs
5,000 LUs simulated by TPNS

VTAM 50K Subarea Release Comparison

	VTAM V4R4			VTAM V4R3		
	Elapsed Time (mm:ss)	CPU seconds (CMC)	CPU seconds (Data)	Elapsed Time (mm:ss)	CPU seconds (CMC)	CPU seconds (Data)
Activation	02:43	146.9	*	02:52	167.8	*
Logon	08:30	523.3	267.2	09:32	592.0	325.0
Failure/React	02:03	34.7	51.2	02:23	44.7	58.7
Re-logon	02:07	47.7	53.5	02:11	52.1	62.7

- **Conclusion:**
  - VTAM V4R4 reduces elapsed and CPU time in all processes as compared to VTAM V4R3

VTAM V4R4 Storage Tuning

- Use worksheets in VTAM Installation and Migration Guide (GC31-8367) - VTAM V4R4 has no Storage Estimates diskette
- Evaluate ECSA value (second value of CSA parm in IEASYSxx) based on following:
  - 1.6Mbytes of VTAM ECSA modules
  - 1.8K ECSA per MNPS session on application host
  - 3.6 ECSA per MNPS session during recovery
  - 3K of default buffers for new pools, CRA4 and CRA8
- TIBUF uses 120K of ECSA - new buffers holding HPR headers for MPC+ data in CSM storage (IOBUF not used at all with MPC+)
  - Code TIBUF=0 if not using MPC+
- Size of BSBUF increased from 248 bytes to 298 bytes
  - Monitor with D NET, BFRUSE
- Average number of IOBUFs on VTAM nodes serving as RTP endpoints is greater than if they support pure APPN traffic
  - Consider increasing BASENO
  - Monitor with D NET,BFRUSE

VTAM V4R4 Throughput Tuning

- *Consider coding MPCLEVEL=NOHPDT on TRL if connection primarily used for small blocksizes*
  - *Not possible to do on some connections -- i.e. VTAM OSA-2 ATM and VTAM to 2216 MPC+ connections*
- *Tune IOBUFsize with HPR (non-MPC+) in mind*
  - *VTAM prefers not to segment IOBUFs. Example:*

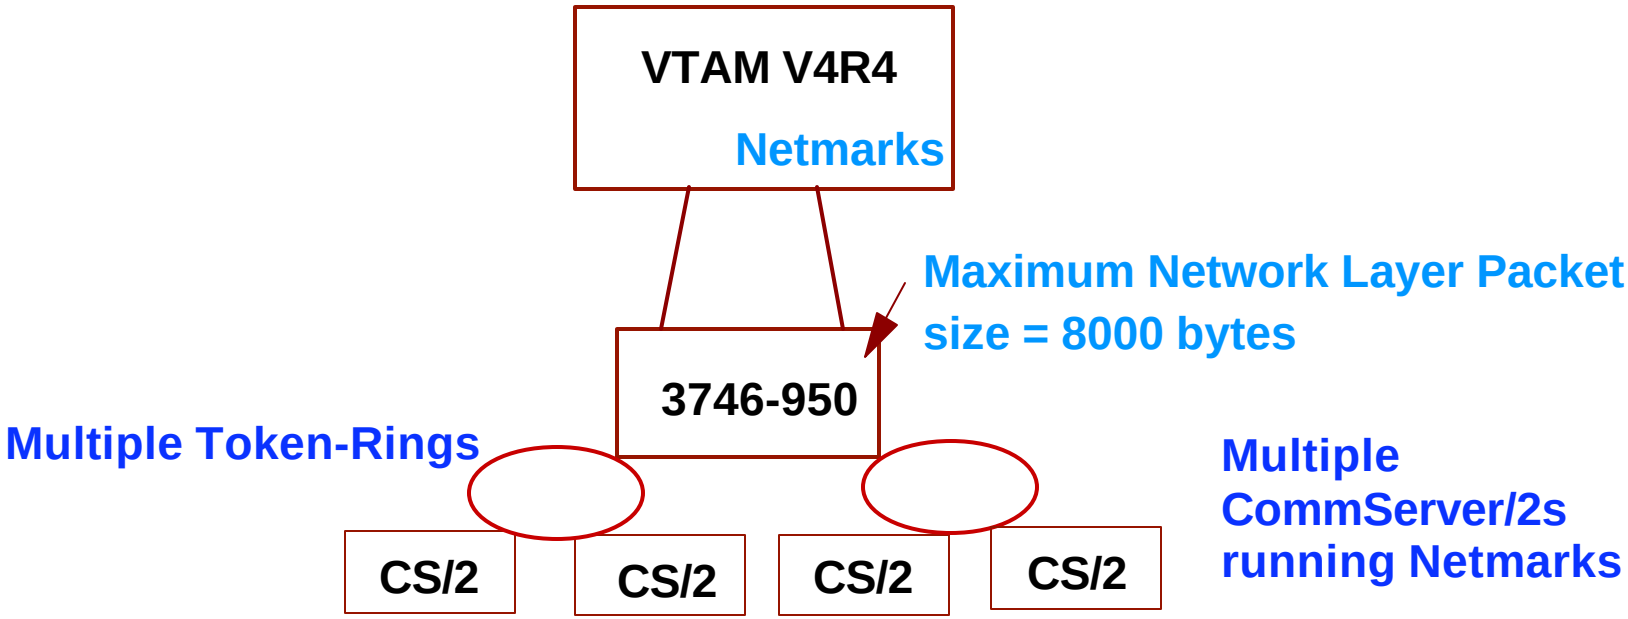

With 16K RUs and IOBUF=3992, VTAM not able to put 2 IOBUFs + HPR header into a single packet fitting within 8000 byte NLP size. Reducing IOBUF to 3900 bytes allows VTAM to send two IOBUFs worth of data in a single packet

Other VTAM V4R4 Performance Items

- *Large VTAM IOBUFsize can significantly improve performance with large blocksize data*
- *MPC+ doesn't use IObuffers -- puts data in 2K blocks of CSM storage and sends maximum allowed on connection at 2K boundary*
- *DELAY default changed to 0 for CTC and NCP connections*
  - ***For CTC default was .1 second***
  - ***For NCP and NCP Data Host was .2 second***
- *SEARCHINFO command allows DISPLAY of outstanding search requests*
  - ***Includes status, DLU node name, etc..***
- *VTAM V4R4 implements adaptive session pacing adaptive window sizes as a receiver*
  - ***In prior releases adaptive pacing window size was constant on the stage where VTAM was the receiver***

Other Interesting VTAM V4R4 Facts

D NET,SESSIONS,SID= shows current pacing values

```
D NET,SESSIONS,SID=FD87C05900BCFE62
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = SESSIONS 124
IST879I PLU/OLU REAL = CSSNET.NETM001    ALIAS = ***NA***
IST879I SLU/DLU REAL = CSSNET.EN001      ALIAS = ***NA***
IST880I SETUP STATUS = ACTIV
IST875I ADJSSCP TOWARDS SLU = ISTAPNCP
IST875I ALSNAME TOWARDS SLU = PU75072
IST933I LOGMODE=#BAT2K  , COS=*BLANK*
IST875I APPNCOS TOWARDS SLU = #BATCH
IST1635I PLU HSCB TYPE: FMCB LOCATED AT ADDRESS X'05AB7C90'
IST1635I SLU HSCB TYPE: BSB LOCATED AT ADDRESS X'05CD4C90'
IST1636I PACING STAGE(S) AND VALUES:
IST1644I PLU--STAGE 1-----|-----STAGE 2--SLU
IST1638I STAGE1: PRIMARY TO SECONDARY DIRECTION - ADAPTIVE
IST1639I        PRIMARY SEND: CURRENT   =     4     NEXT =     7
IST1640I        SECONDARY RECEIVE       =     7
IST1641I STAGE1: SECONDARY TO PRIMARY DIRECTION - ADAPTIVE
IST1642I        SECONDARY SEND: CURRENT =    54     NEXT =    63
IST1643I        PRIMARY RECEIVE         =    63
IST1638I STAGE2: PRIMARY TO SECONDARY DIRECTION - ADAPTIVE
IST1639I        PRIMARY SEND: CURRENT   =    50     NEXT =    64
IST1641I STAGE2: SECONDARY TO PRIMARY DIRECTION - ADAPTIVE
IST1643I        PRIMARY RECEIVE         =    63
IST314I END
```

APPN/HPR Performance

- HPR can significantly improve performance
- RTP endpoints implement:
  - ARB - allowing data to flow into network at high rate without causing congestion
  - Selective frame retransmission
  - Non-disruptive session switching
- Intermediate nodes (ANR nodes) can:
  - Increase throughput due to faster routing mechanism than with APPN ISR
  - Support more traffic due to reduction in processor and storage requirements
- RTP endpoints may require somewhat increased processor and storage requirements due to end to end session recovery and flow control
- Recent HPR performance information follows

APPN/HPR VTAM V4R3 Channel Measurement

- *Measurement used to compare HPR with APPN for session endpoints:*
  - *Throughput*
  - *Storage*
  - *CPU Cycles*
- *4 Mbyte file transferred repeatedly from client to server*
  - *Data in storage*
  - *Large IOBUFs*
  - *Large pacing windows*
  - *No coattailing delays*

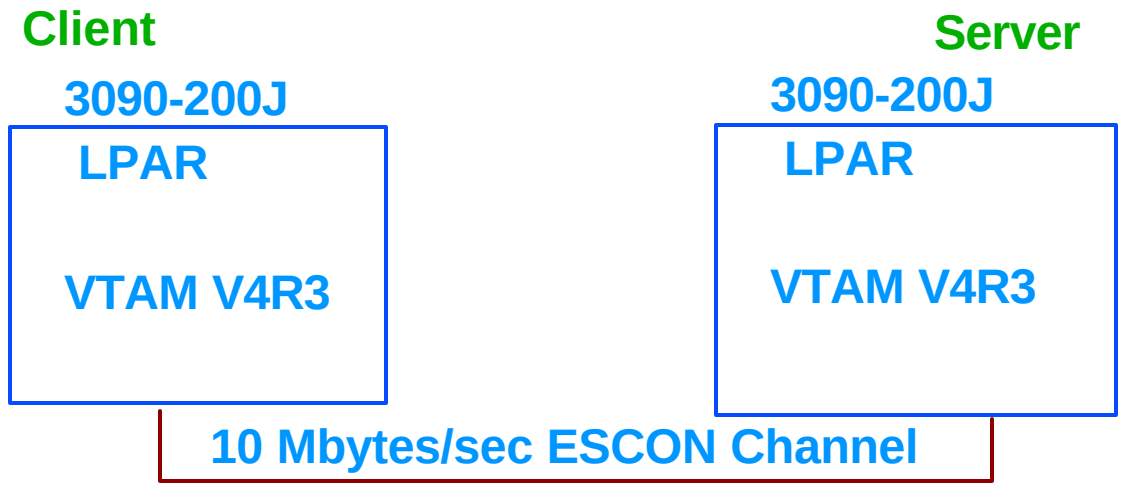

VTAM V4R3 HPR/APPN Measurement Results

	HPR	APPN
Thruput (Mbytes/Sec)	8.23	8.36
Client Host		
Storage		
IOBUF (Max Used)	2191	1702
CSA	11177K	11144K
Private	6360K	6344K
Utilization		
CP0	35.4%	33.5%
CP1	37.4%	29.3%
Server Host		
Storage		
IOBUF (Max Used)	313	80
CSA	11079K	11052K
Private	6100K	6100K
Utilization		
CP0	66.7%	65%
CP1	48.9%	45.5%

17MBytes/sec Channel VTAM V4R3 Measurements

- *Measurement used to compare 17 Megabyte ESCON channel for VTAM-VTAM communication with 10 Megabyte ESCON channel*

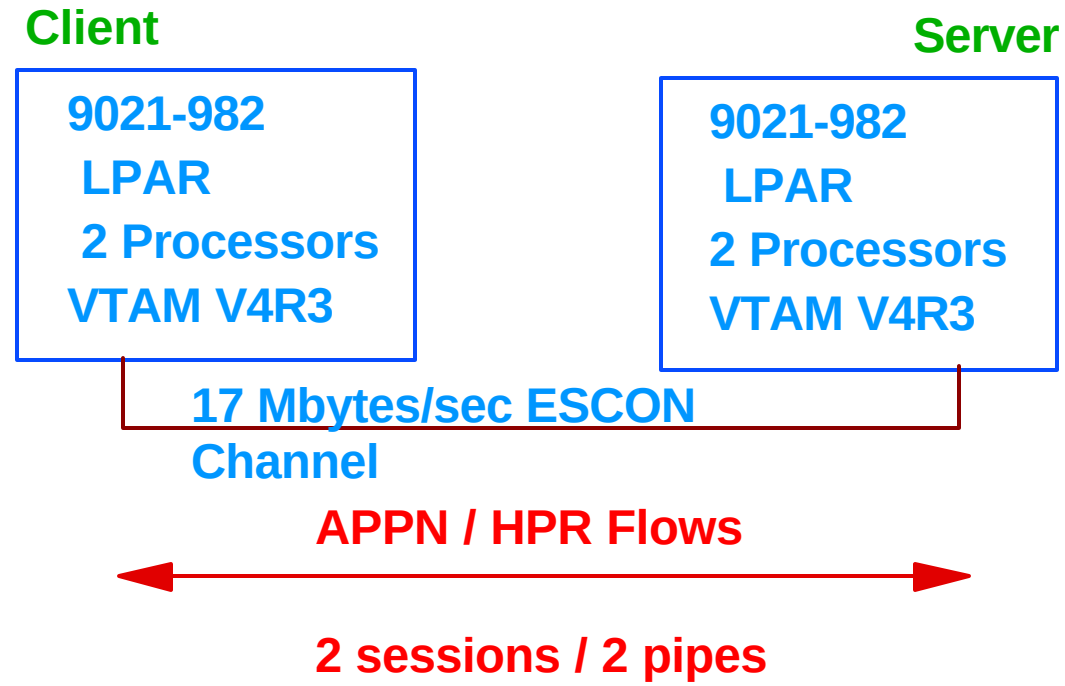

- *4 Mbyte file transferred repeatedly from client to server*
  - *Data in storage*
  - *Large IOBUFs*
  - *Large pacing windows*
  - *No coattailing delays*

17 Mbytes/Sec Channel Measurement Results

	HPR	APPN
Thruput (Mbytes/Sec)	14.76	15.09
Client Host Utilization	22%	20%
Server Host Utilization	38%	34%

VTAM V4R3/NCP V7R3 HPR/APPN Measurement

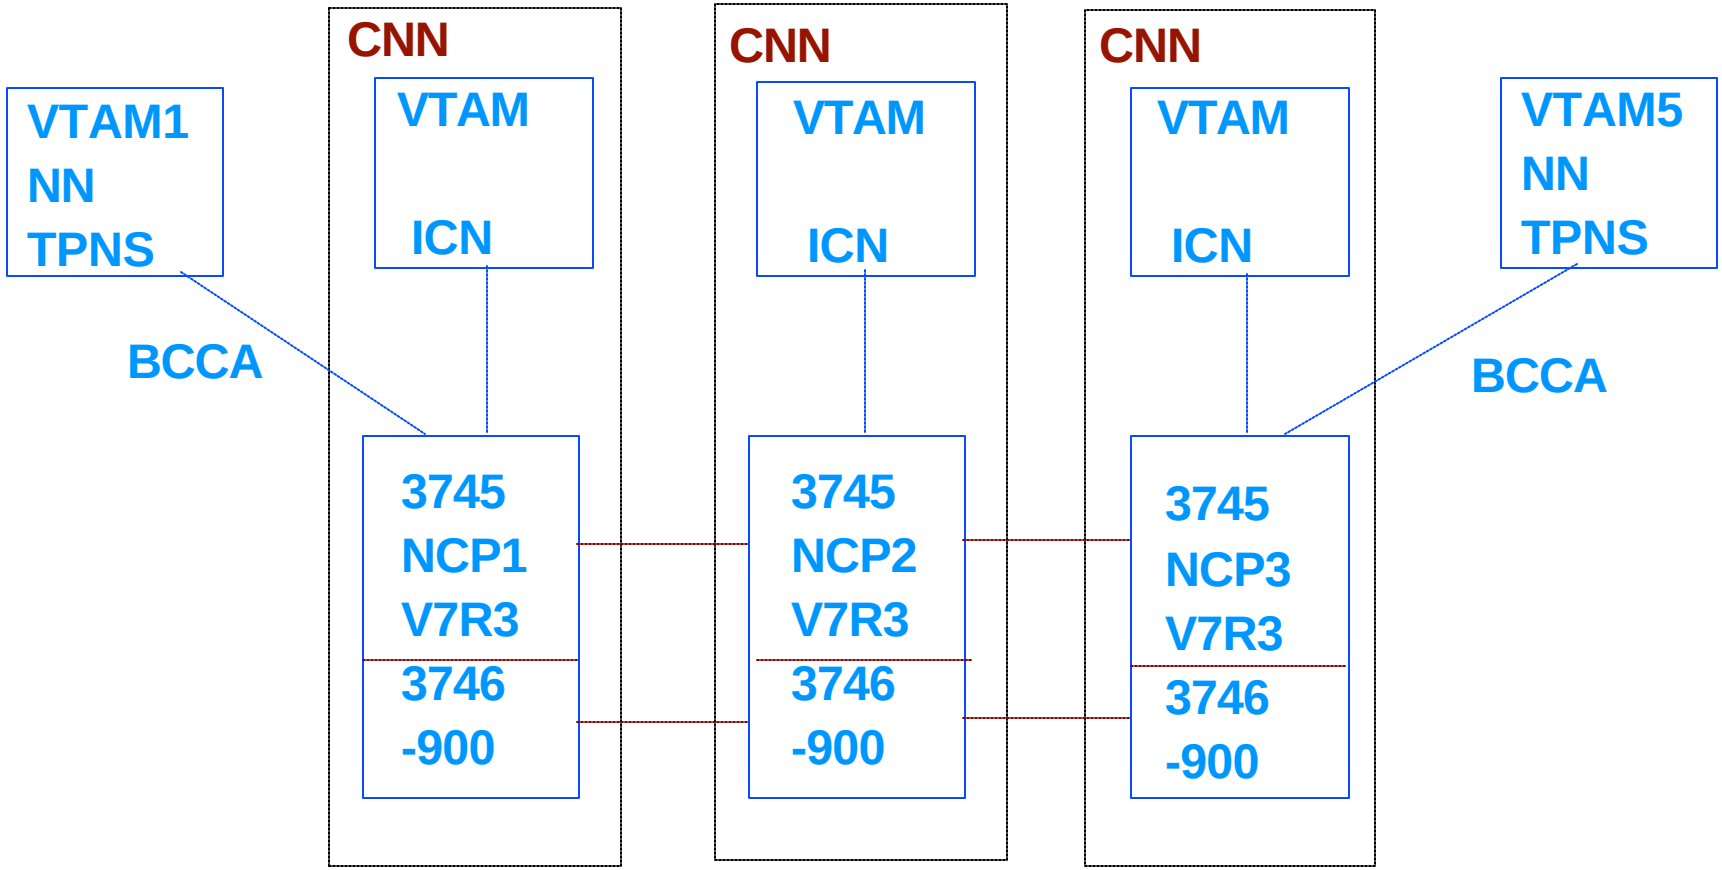

- *TPNS used to simulate interactive workload*
- *200 sessions/200 pipes - 128/128 bytes per transaction*
- *APPN vs. HPR compared for this configuration/workload*

VTAM/NCP HPR Measurement Results

	HPR	APPN
Thruput Trans/Min	21,140	17,930
Host Utilization Storage (Client/Server) IOBUF CSA Private	37.5% (4.4% increase for same thruput) 430K / 498K 3201K / 3258K 8485K / 7794K	28.1% 119K / 97K 2835K / 2740K 5281K / 5465K
NCP Utilization NCP1 NCP2 NCP3 Storage Routing Capacity at same utilization	 57.6% 32.9% 41.4% Savings of 400 bytes per NCP per session 2-3X (Frame Relay on 900 could be greater)	 95.0% 65.5% 93.6% X

NCP V7R3 Comparison (CF3745)

Subarea:

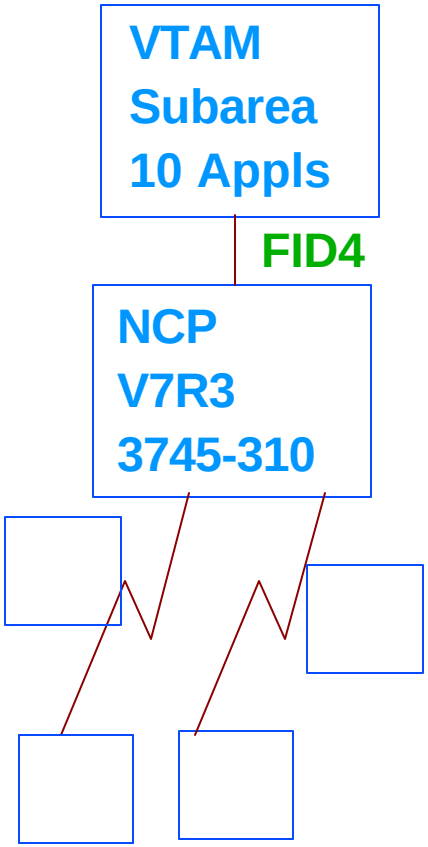

50 SDLC 56K links
10 PUs/Line
10 LUs/PU
5000 Sessions

APPN:

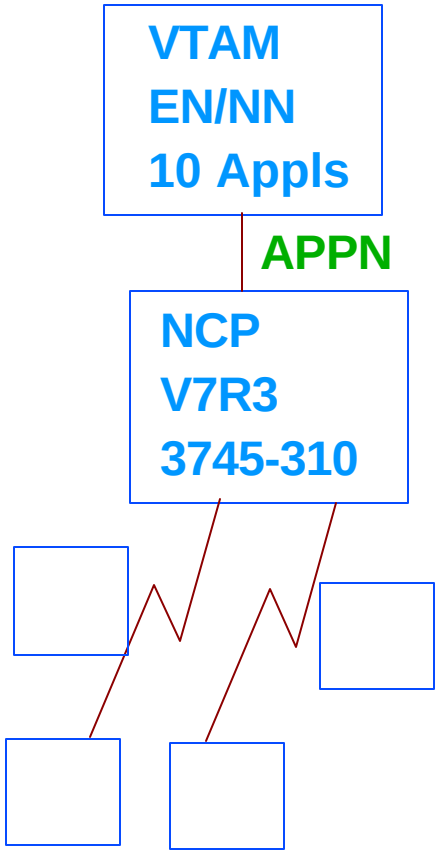

50 SDLC 56K links
10 PUs/Line
10 LUs/PU
5000 Sessions

HPR:

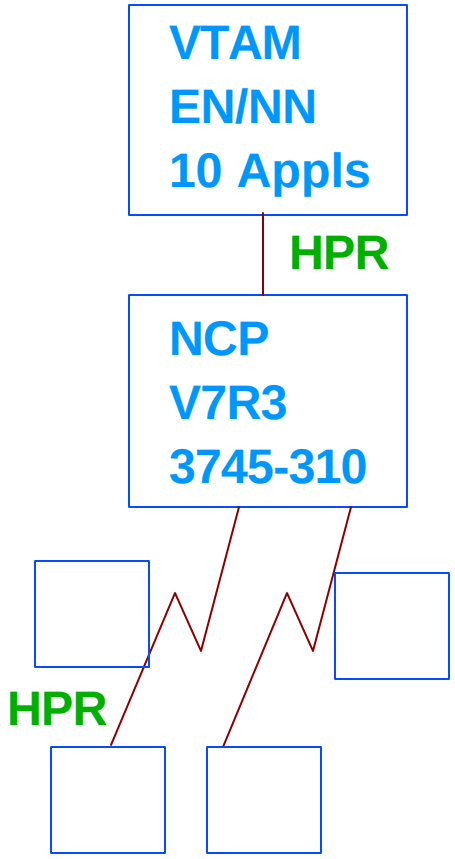

50 SDLC 56K links
10 HPR Nodes/Line
5000 Sessions

- *Also modeled using Frame Relay links - non-ERP for HPR*

NCP V7R3 CF3745 Comparison Results

SDLC Configurations:

	Subarea	APPN	HPR
CCU Utilization	78%	87%	57%
Channel Utilization	35%	34%	37%
SDLC Line Utilization	70%	70%	68%
Storage	4.19Mbytes	5.31Mbytes	2.45Mbytes

Frame Relay Configurations:

	Subarea	APPN	HPR
CCU Utilization	80%	90%	49%
Channel Utilization	35%	34%	38%
Frame Relay Line Utilization	68%	68%	66%
Storage	4.55Mbytes	5.68Mbytes	3.07Mbytes

1997 Benchmark ISR/HPR Comparisons

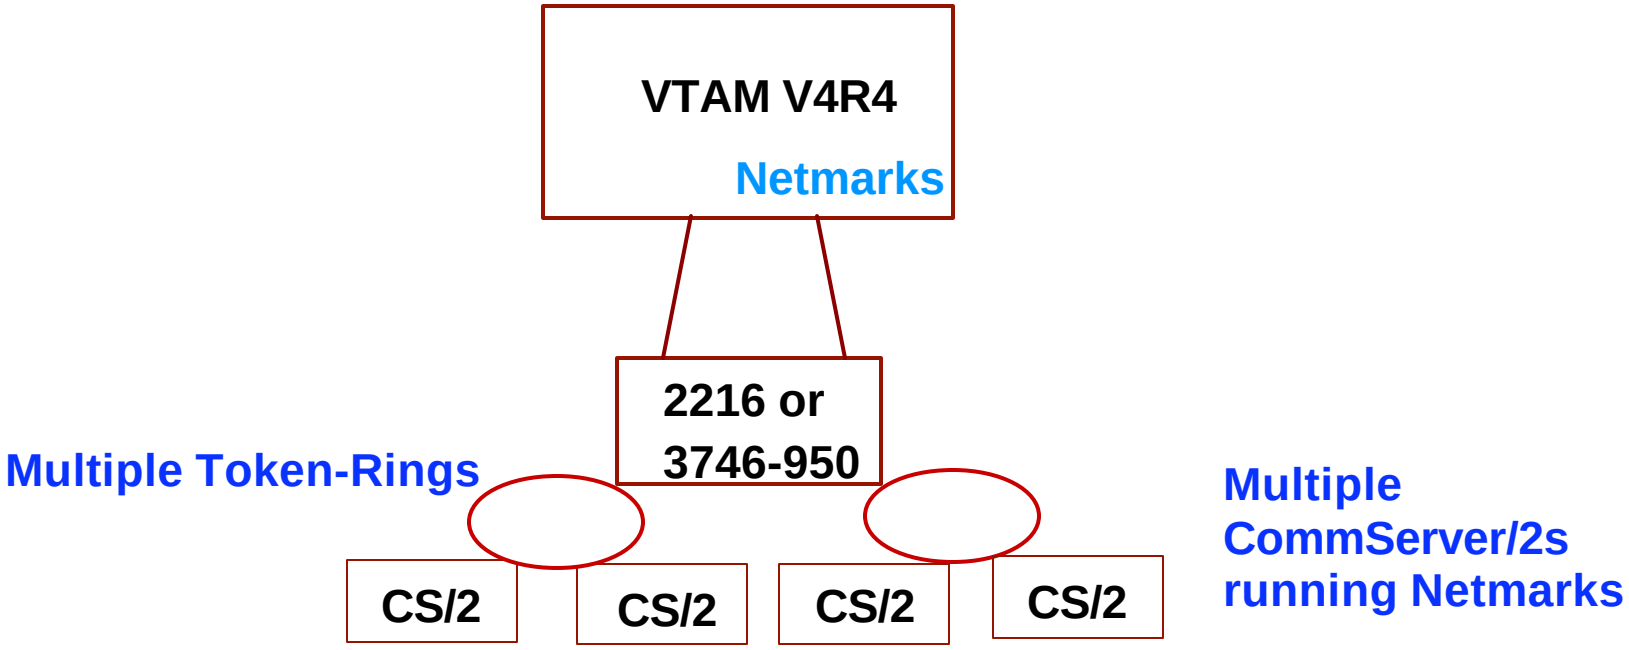

- *3746-950 HPR and ISR Configuration Results*
  - *10-30% throughput increase with HPR versus ISR for large blocks*
  - *Minimal VTAM cycle increase for RTP endpoint function*
- *2216 HPR/ISR*
  - *Significant throughput (3-4 times) increase with HPR versus ISR for large blocksizes*
  - *VTAM cycle requirements for RTP endpoint much less than APPN ISR due to LSA channel protocol usage with ISR configuration*

Ganymede Software Test Report

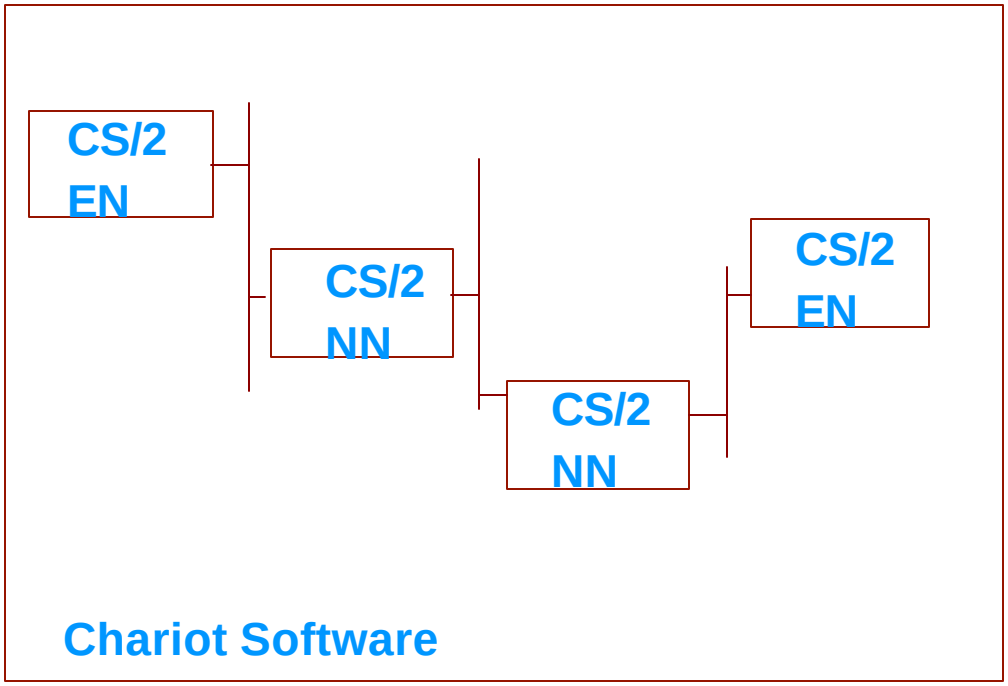

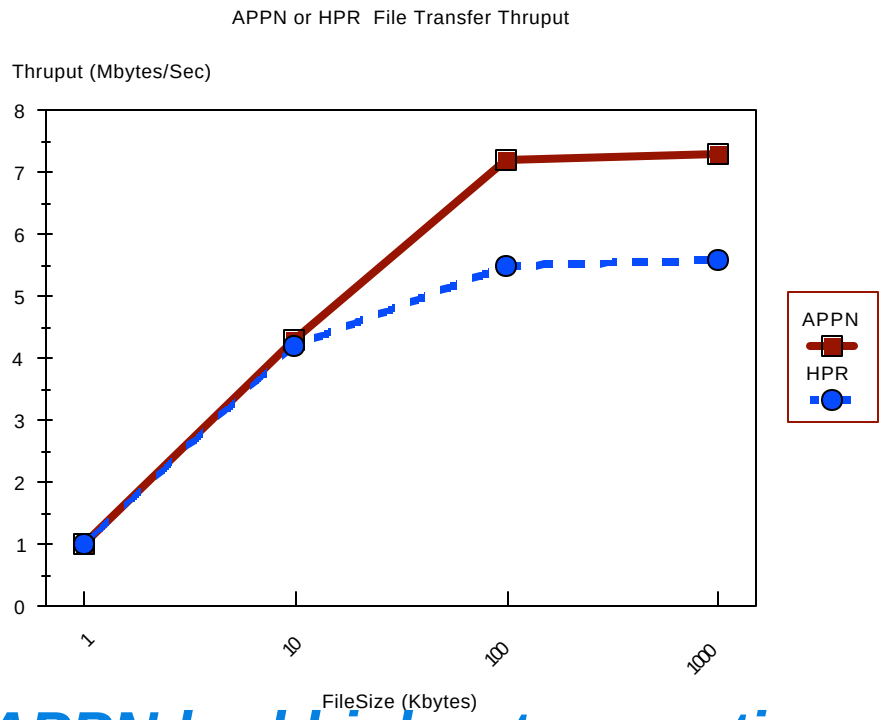

- *In Ganymede Test configuration, APPN had higher transaction rate and better throughput with the larger file sizes than HPR*
- *End node processors were highly utilized, RTP function added 10-30% more processor utilization*
- *With HPR, NN utilization stayed at about 25%, while grew with load in APPN*
- *To understand HPR's effects, test with planned configurations*
  - ***ANR nodes should support much more traffic but don't always increase throughput rate -- different platforms perform differently***

3172 (LSA) Revised Formula

- **Formula provided in 1993 to estimate additional CPU time needed for 3172 XCA versus NCP supported resources**
  - **Original formula was low because assumed traffic was INN only**
  - **Formula has been enhanced to estimate BNN functions**
  - **Formula has been enhanced to estimate routing functions**
- **Formula for Subarea, but roughly applies to APPN**

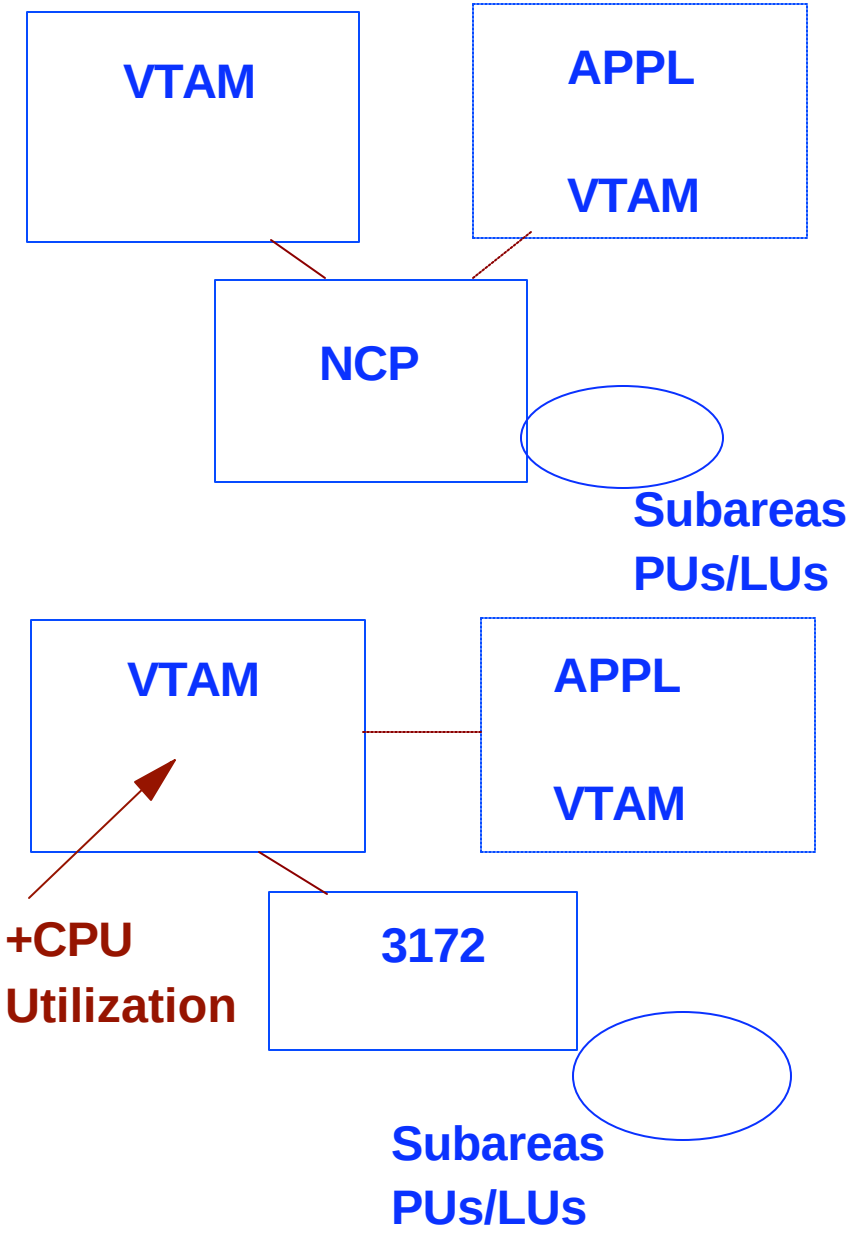

Revised Formula for 3172 VTAM Cycles

The formula is:

Compute the increase in host utilization due to OUTBOUND traffic:

Uo = Lo * (1706 + 26*No)/M * 100% (INN)

Uo = Lo * (1706 + 1585 + 26*No)/M * 100% (BNN)

Compute the increase in host utilization due to INBOUND traffic:

Ui = Li * (2116 + 7*Ni)/M * 100% (INN)

Ui = Li * (2116 + 763 + 7*Ni)/M * 100% (BNN)

Compute the increase in host utilization due to SESSION ROUTED traffic:

Ur = Lr * (2116 + 1706 + 3500 + 26*no + 7*Ni)/M * 100% (Subarea)

Ur = Lr * (2116 + 1706 + 5800 + 26*no + 7*Ni)/M * 100% (APPN)

Add the results of each computation

Notation:

Lo Average Outbound PIUs/second

Li Average Inbound PIUs/second

Lr Average routed PIUs/second (INN or APPN ISR)

M Host MIPs in Instructions/sec, i.e. 20MIPs = 20,000,000

No Average IOBUFs/outbound PIU

Ni Average IOBUFs/inbound PIU

3172 VTAM CPU Formula Example

Compute increase in VTAM CPU utilization for 3172 BNN environment::

IOBUFsize = 512, Host MIPs = 20

300 Inbound 300 byte PIUs/sec, 400 Outbound 1200 bytes PIUs/sec

Uo = 400 * (1706 + 1585 + 26 * 3)/20000000 * 100% = 6.74%

Ui = 300 * (2116 + 763 + 7 * 1)/20000000 * 100% = 4.34%

Increase in processor utilization for VTAM cycles = 11.08%

Compute additional increase for routing 100 PIUs/sec of 50 bytes to an APPN channel attached node:

Ur = 100 * (2116+1706+5800+26+7)/20000000 * 100 = 4.83%

The total increase in CPU utilization for VTAM cycles is 15.91%, e.g. a processor averaging 39.5% busy with VTAM cycles would increase to 55.41%

1996 Benchmark with Cisco,NCP and 3746

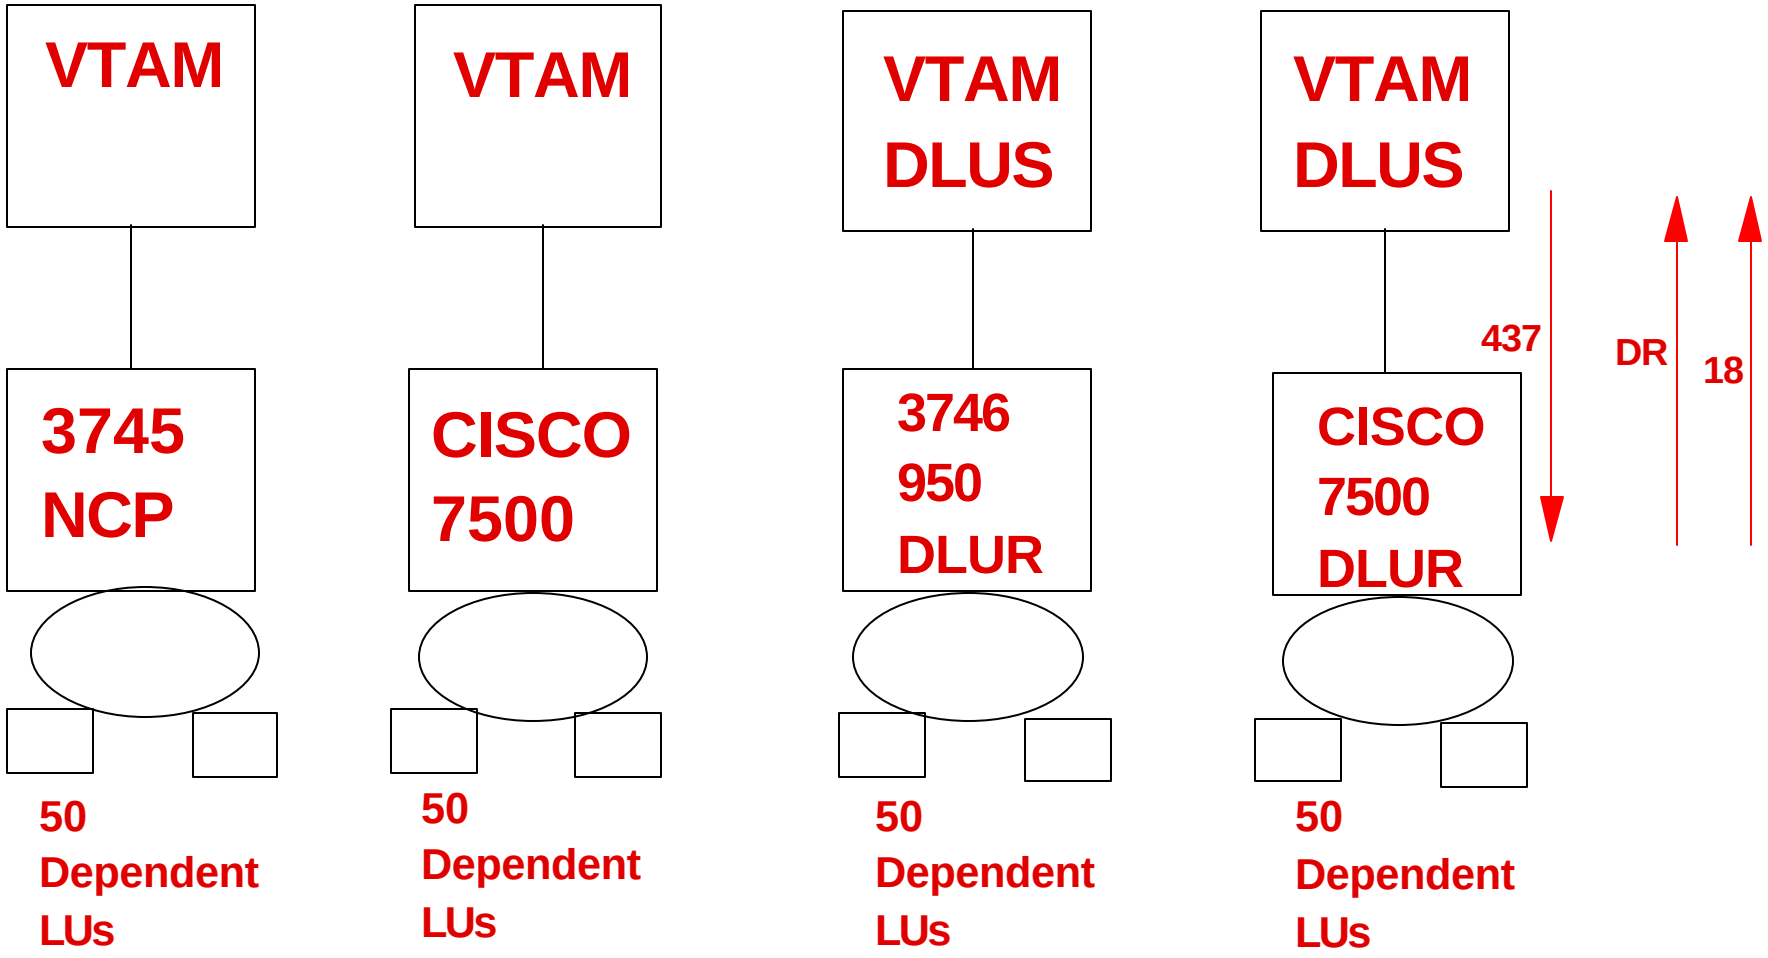

- *Equalized Performance Comparison:*
  - *3745 subarea to APPN causes 2.92% increase in VTAM cycles*
  - *3745 subarea to CISCO subarea causes 46.9% increase in VTAM cycles*
  - *3745 subarea to CISCO APPN cause 49.8% increase in VTAM cycles*
  - *CISCO APPN instead of 950 APPN cause 45.6% increase in VTAM cycles*

1997 Benchmark VTAM Cycles Comparisons

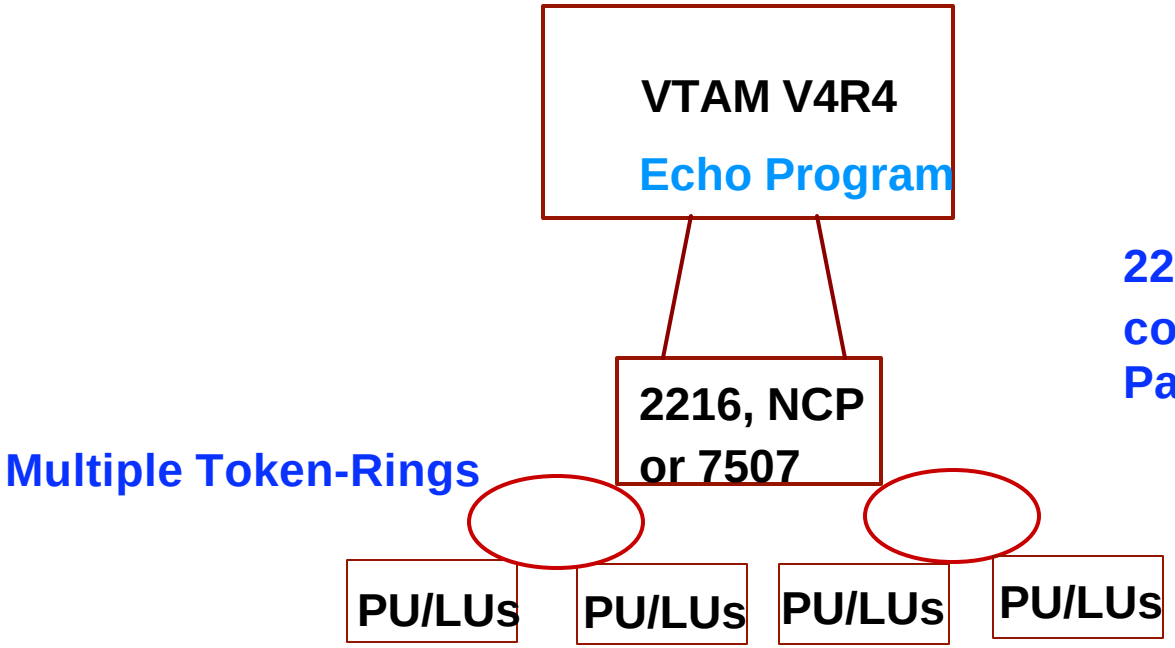

2216 and 7507 configured in Passthrough Mode

TPNS simulating about 1000 LUs, 4 per PU

S/390 Utilization % (VTAM & MVS)

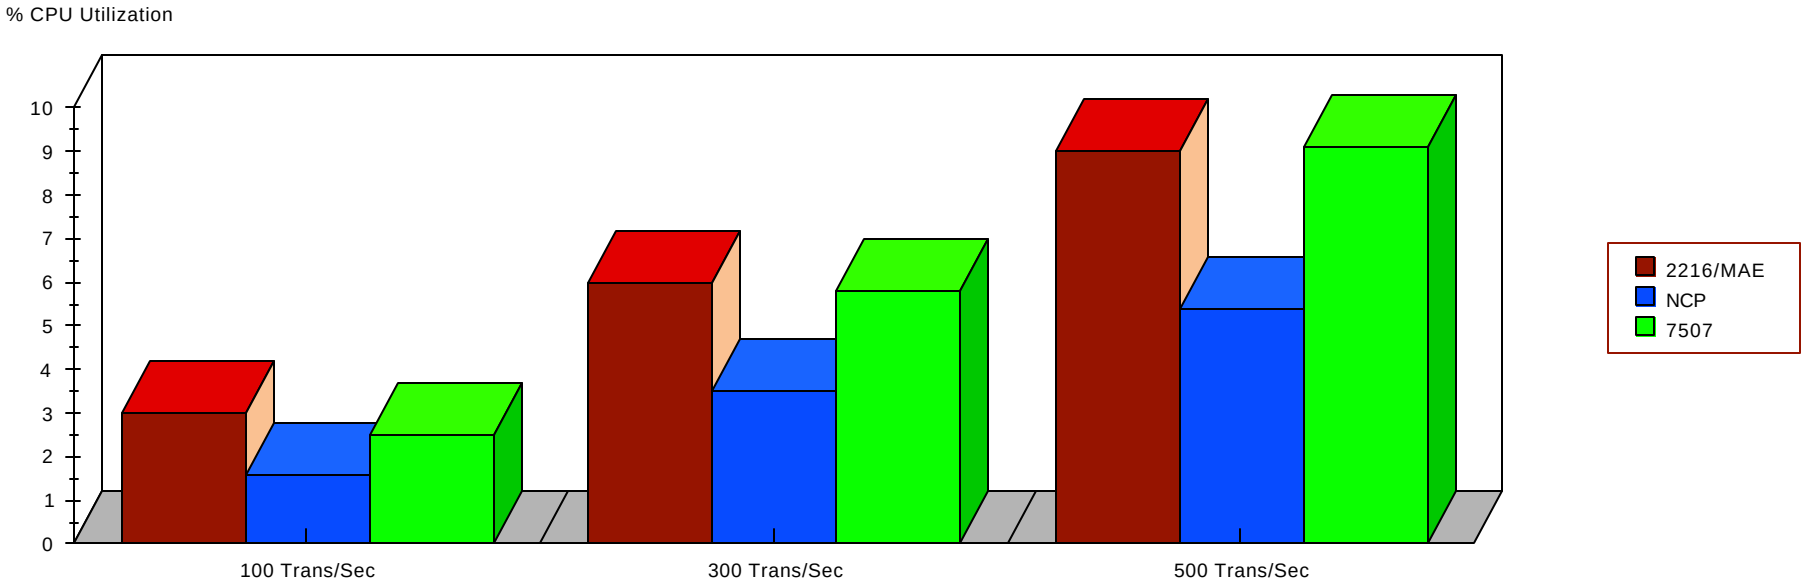

1997 Benchmark VTAM Cycles Comparisons

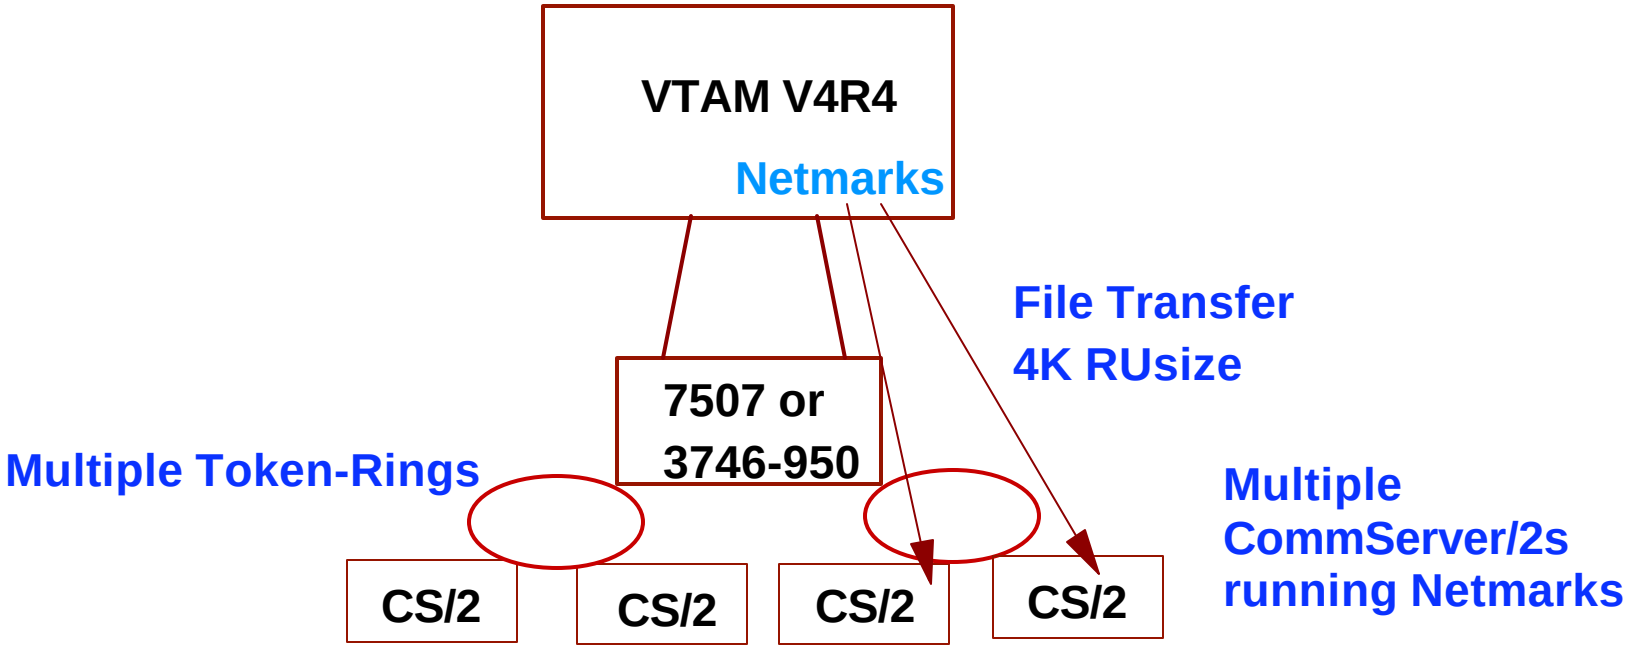

- *3746-950 and 7507 running as APPN NNs*
- *Both configurations running APPN ISR*
- *7507 uses LSA channel protocol, 3746-950 uses CDLC channel protocol*
  - *7507 required 65% more S/390 processor time per byte transferred*
  - *3746-950 required significantly less S/390 processor time per byte transferred with 8K and 16K than with 4K or 2K RUsizes*

Other VTAM / NCP Performance Highlights

- **APPN Multiple Network Connectivity (Border Node)**
  - **10K LU XNet logon took 75% less elapsed time compared to SNI configuration**
  - **VTAM V4R4 with or without NCP V7R5 supports HPR across Border Node connections**
    - NCP HPR Border Node significantly reduces processing and storage as compared to SNI
- **Virtual Route Transmission Groups**
  - **10K LU Xdomain logon 40% faster in elapsed time and used 10% less CPU cycles than APPN configuration**
- **Results of VTAM V4R3 and V4R2 benchmarks documented in ASKQ WSC Flashes**
  - **Use search argument "VTAM Performance Benchmark"**

Summary

- *VTAM V4R4 Performance Information*
  - *VTAM V4R4 Benchmark Results*
  - *VTAM V4R4 performance related enhancements*
  - *Storage and Cycle Tuning general guidelines*
- *HPR Performance Information*
  - *Review of HPR performance Information*
  - *Recent HPR benchmark performance information*
  - *HPR tuning general guidelines*