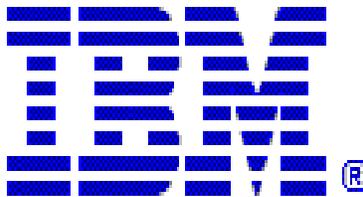


# **IBM Communication Controller for Linux on System z9 and zSeries V1R2 (CCL)**

—

## **System z9 and zSeries CPU Capacity Planning Information for SNI and Boundary Function Workload**



Copyright 2006 IBM Corporation

Bob Perrone – [bperrone@us.ibm.com](mailto:bperrone@us.ibm.com)  
Alfred B Christensen – [alfredch@us.ibm.com](mailto:alfredch@us.ibm.com)

## Table of contents

Disclaimer .....	3
CCL release and maintenance levels.....	3
Summary of results .....	3
Workload characteristics .....	4
SNI transactional workload .....	4
Boundary function transactional workload .....	5
Test environment overview.....	5
Architectural differences between IBM 3745/46 and CCL.....	5
IBM 3745/46 SNI test environment.....	6
CCL SNI test environment.....	6
IBM 3745/46 boundary function test environment.....	8
CCL boundary function test environment .....	8
CCL memory requirements.....	10
SNI transactional workload - results.....	11
Estimating SNI CCL CPU utilization .....	13
Boundary function transactional workload - results.....	17
Estimating boundary function CCL CPU utilization .....	19
Quick CCL V1R2 CPU utilization conversion tables .....	23
SNI from IBM 3745-31A with TIC2 to z9-109 .....	23
SNI from IBM 3745-31A with TIC3 to z9-109 .....	24
BF from IBM 3745-31A with TIC2 to z9-109.....	24
BF from IBM 3745-31A with TIC3 to z9-109.....	25
Notes .....	25
Samples .....	26
IBM 3745 17A SNI workload to CCL V1R2 on a z990 .....	26
IBM 3745 21A boundary function workload to CCL V1R2 on a System z9 .....	26

### ***Disclaimer***

This document contains System z9 CPU capacity planning guidelines for migrating SNA Network Interconnection (SNI) and boundary function (BF) transactional workload from an IBM 3745/46 environment to a CCL environment. If you are not planning to deploy CCL on a System z9 CPU, please use the usual Large Systems Performance Ratio (LSPR) data to adjust the CPU capacity in this document to the level of zSeries hardware you are planning to use.

***Disclaimer:*** All performance data contained in this document was obtained in the specific operating environment and under the conditions described and is presented as an illustration only. Performance obtained in other operating environments may vary and customers should conduct their own testing.

### ***CCL release and maintenance levels***

The first version of CCL shipped in late March 2005 as CCL V1R1.

A performance APAR for CCL V1R1 (APAR number LI70826) was made available mid August 2005. This document will refer to a CCL V1R1 system with the performance APAR applied as CCL V1R1+.

CCL V1R2 shipped mid November 2005. CCL V1R2 further improves throughput and reduces CPU consumption as compared to CCL V1R1+.

The main focus of this document is to provide system z9 CPU capacity planning data for the CCL V1R2 level.

CCL V1R1 and CCL V1R1+ were both limited to use of OSA LCS copper-based LAN connectivity for SNA network flows – both upstream to VTAM and downstream to other NCPs or boundary function SNA devices.

CCL V1R2 added numerous new connectivity options for SNA network flows. On a System z9, VTAM or TPF can connect to a CCL V1R2 NCP via a shared OSA adapter operating in OSA for NCP (OSN) mode. The OSN connectivity appears to VTAM, TPF, and the NCP as an ESCON channel over which the traditional CDLC channel protocol is used. On zSeries z890 and z990 or on a System z9, CCL V1R2 also adds support for OSA QDIO Layer-2 copper or fiber-based LAN connectivity for SNA network flows. On all supported hardware levels, two CCL V1R2 NCPs communicating with each other (INN or SNI traffic) can exploit a new CCL V1R2 function known as IP Transmission Group (IP-TG) that allows the two CCL environments to exchange the INN/SNI traffic over a TCP connection via an intermediate IP-based network infrastructure. Use of these new CCL V1R2 connectivity options has in some cases significant impact on both throughput and CPU consumption.

### ***Summary of results***

A CCL V1R2 environment running on a System z9 CPU can deliver significantly improved throughput (transactions per second) and response times as compared to a similar NCP-based workload running on a real IBM 3745/46 environment.

The CCL CPU utilization generally scales following a linear model. Given a known set of workload characteristics and transaction rates, the CCL CPU utilization can be roughly estimated using simple linear algorithms as shown in this document.

A CCL V1R2 environment running on a single System z9 IFL engine can support an SNI transactional workload as defined in this document, of over 2200 transactions per second, at

## CCL V1R2 – System z9 and zSeries CPU Capacity Planning

which rate the utilization of the System z9 CPU reaches 86%. This is in some cases between 5 and 6 times more transactions per second than what can be provided by an IBM 3745/46 31A configuration. With CCL V1R2 running on a system z9, it is possible to consolidate SNI workload from up to five IBM 3745 31A configurations (CCU utilization at around 70%) into a single system z9 IFL engine.

OSN connectivity between VTAM and a CCL NCP improves performance about 40% as compared to a shared LAN between VTAM and a CCL NCP. This capacity planning document will highlight the performance differences between shared LAN and OSN for VTAM to CCL NCP connectivity for both SNI and boundary function workload.

IP-TG connectivity between two CCL NCPs also improves performance as compared to SNA LLC2 over a shared LAN between the two CCL NCPs. An SNI test environment using OSN between VTAM and the CCL NCPs and IP-TG between the two CCL NCPs performs about 30% better than a similar environment using SNA LLC2 over a shared LAN between the two NCPs. This capacity planning document will highlight the performance differences between SNA LLC2 over a shared LAN and IP-TG for CCL NCP to CCL NCP connectivity for SNI workload.

Use of QDIO Layer-2 for SNA LLC2 traffic between a CCL NCP and a LAN does not appear to have performance benefit over other LAN connectivity. However, use of QDIO Layer-2 provides much improved OSA port sharing capabilities and the opportunity to transport SNA LLC2 traffic over fiber-optic LAN ports, such as Gigabit and 10 Gigabit OSA ports.

When comparing CCL CPU utilization to IBM 3745 CCU utilization, many factors influence the comparison. The most significant factor appears to be the hardware configuration of the IBM 3745/46 environment – in particular whether TIC2 or TIC3 adapters are used for LAN connectivity to the IBM 3745/46. An attempt to estimate CCL CPU utilization based on existing IBM 3745 CCU utilization should not be made without a clear understanding of the current IBM 3745/46 hardware configuration.

### ***Workload characteristics***

This document covers transactional SNI workload and transactional NCP boundary function workload.

### **SNI transactional workload**

The transactional SNI workload that was used in the performance tests is intended to simulate interactive transactions between two SNI partners, where one SNI partner sends short transaction requests to the other, and the second SNI partner responds with slightly longer messages. The workload resembles a traditional 3270 terminal dialog or typical LU 0 and LU 6.2 transactions being exchanged between the two business partners.

The transactional workload had the following characteristics:

- SNA APPC sessions between two z/OS systems were used to generate the transactional workload.
- Each transaction used 100 bytes in and 800 bytes out with a 330 millisecond think time between transactions.
- The workload was generated using the IBM Application Workload Modeler (AWM) tool. You can get information about AWM at <http://www.ibm.com/software/network/awm/>.
- A series of runs were performed where the number of APPC sessions was gradually increased from 25 to 50, then to 100, 150, and finally 175 sessions.
- VTAM IOBUF size was set to 932 in all test runs.
- NCP MAXBFRU was set to 20 in all test runs.
- CCL NCP MAXOUT was set to 7 for CCL to VTAM.

### Boundary function transactional workload

The BF transactional workload had the following characteristics:

- SNA APPC sessions were used to generate the BF transactional workload.
- Each transaction used 100 bytes in and 800 bytes out with a 330 millisecond think time between transactions
- The workload was generated using the IBM Application Workload Modeler (AWM) tool. You can get information about AWM at <http://www.ibm.com/software/network/awm/>.
- A series of runs were performed where the number of APPC sessions was gradually increased from 25 to 50, then to 100, 150, and finally 180 sessions.
- VTAM IOBUF size was set to 932 in all test runs.
- NCP MAXBFRU was set to 20 in all test runs.
- CCL NCP MAXOUT was set to 7 for CCL to VTAM.

### Test environment overview

For the SNI performance tests that are the base for capacity planning guidelines in this document, two z/OS systems in two separate SNA networks (NETA and NETB) and two NCPs were configured in an SNI configuration using a null net (NETX). The transactional workload was generated using IBM Application Workload Modeler (AWM), where both the AWM sender component and the AWM responder component were deployed on z/OS systems.

For the NCP boundary function performance tests, a single z/OS system was used. The peripheral nodes were simulated using IBM Application Workload Modeler (AWM) with all peripheral nodes attached through a LAN.

The NCP level used in all measurements was NCP V7R8.1.

The Linux images that were used in the performance tests were all running SUSE Linux SLES9 with SP2 (a Linux 2.6 kernel) and IBM Communication Controller for Linux on System z9 and zSeries. Measurements were made with the original CCL V1R1 code, the CCL V1R1+ code (with the August 2005 performance APAR applied), and with the CCL V1R2 code.

The Linux images all ran as guests under z/VM 5.1 and each had one virtual CP and 512 MB memory assigned.

Please note that the CCL V1R1 and V1R1+ test runs were done with CCL running on a z990 processor. The CPU utilization numbers were then converted to System z9 utilization to enable comparison with the CCL V1R2 data. When CCL V1R1 and V1R1+ both appear to reach maximum throughput at around 74% utilization of a System z9, then remember that 74% utilization on a System z9 is approximately equal to 100% utilization of a z990. Both CCL V1R1 and CCL V1R1+ running on a z990 were limited by the CPU capacity in that environment.

### Architectural differences between IBM 3745/46 and CCL

An IBM 3745/46 environment can be a very complex environment that may consist of multiple specialized processors that assist the general IBM 3745 CCU in processing the NCP workload. Examples of specialized processors are the ESCON channel processor that handles low-level ESCON channel processing and the token-ring processor that handles the link-layer control functions for the token-ring LANs that are attached through TIC3 adapters. The type, the number, and the mix of assist processors in an IBM 3745/46 configuration is an installation choice and varies significantly. The use of assist processors offloads various types of processing from the general IBM 3745 CCU to the assist processors and therefore reduces the IBM 3745 CCU utilization. Comparing CCU utilization between two IBM 3745/46 environments can be quite complex, since the two environments may support comparable workloads, but each shows a different CCU utilization due to different sets of assist processors. An environment with many

## CCL V1R2 – System z9 and zSeries CPU Capacity Planning

combined processors will in general not show a linear scalability, but will tend to show a gradually decreasing throughput as workload increases.

The CCL environment is best compared to an IBM 3745 environment without any assist processors, which means that the System z9 or zSeries CPU performs general NCP functions as well as all low-level link-layer control functions. From a structural perspective, the CCL environment resembles an IBM 3745 using TIC2 token-ring adapters. The CCL environment will therefore also in general exhibit a more linear scalability than an IBM 3745/46 environment.

### IBM 3745/46 SNI test environment

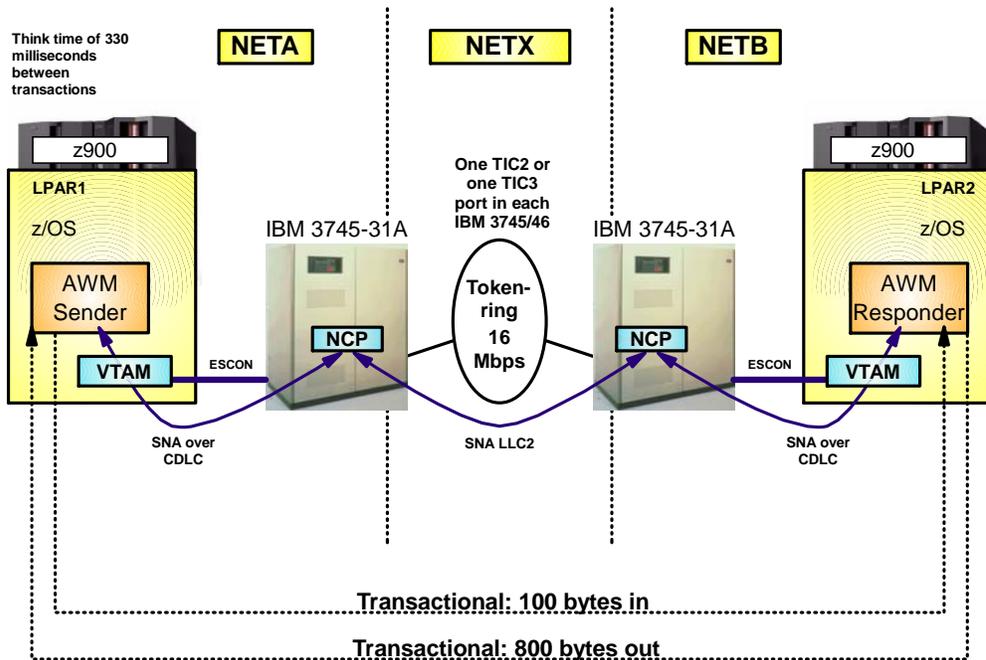


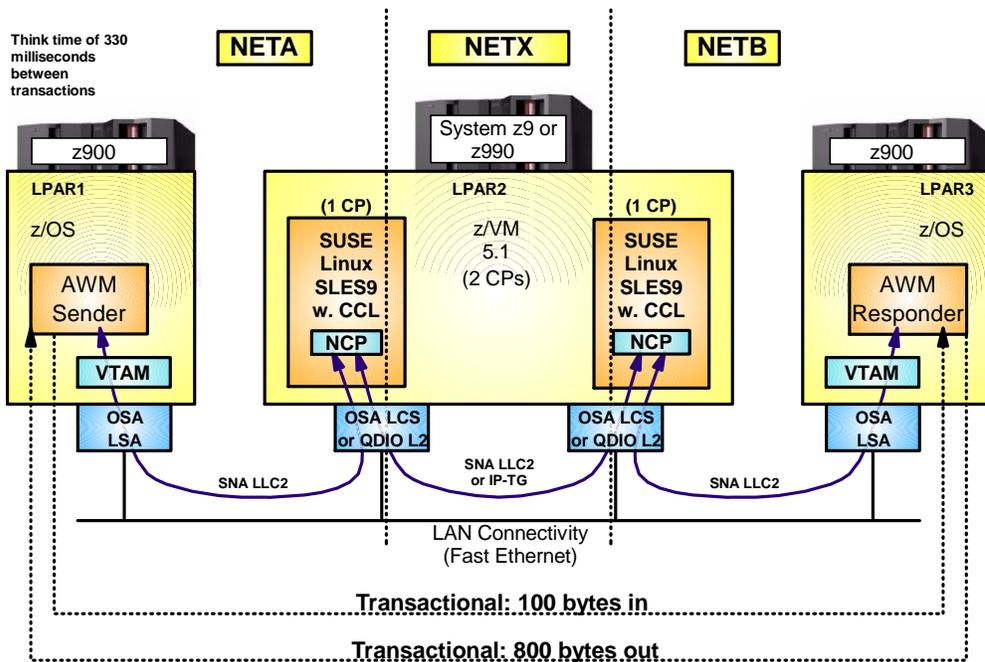
Figure 1: IBM 3745/46 SNI test environment

The two z/OS systems were ESCON channel-attached to each of their IBM 3745-31A. The two IBM 3745-31As were connected together over a 16 Mbps token-ring network using TIC2 or TIC3 adapters. One set of test runs were done using TIC2 adapters and another set of runs were done using TIC3 adapters. The TIC2 token-ring adapters are located in the IBM 3745 frame and link-layer control functions are performed by the NCP code. The TIC3 token-ring adapters are located in the IBM 3746-900 module under control of a Token-Ring Processor (TRP) to which the NCP offloads the link-layer control functions.

### CCL SNI test environment

There were two different CCL test environments used for the performance tests for this document. One scenario used shared LAN between VTAM and the CCL – VTAM using an OSA LSA port to a Fast Ethernet, and CCL using an OSA LCS port to the same Fast Ethernet.

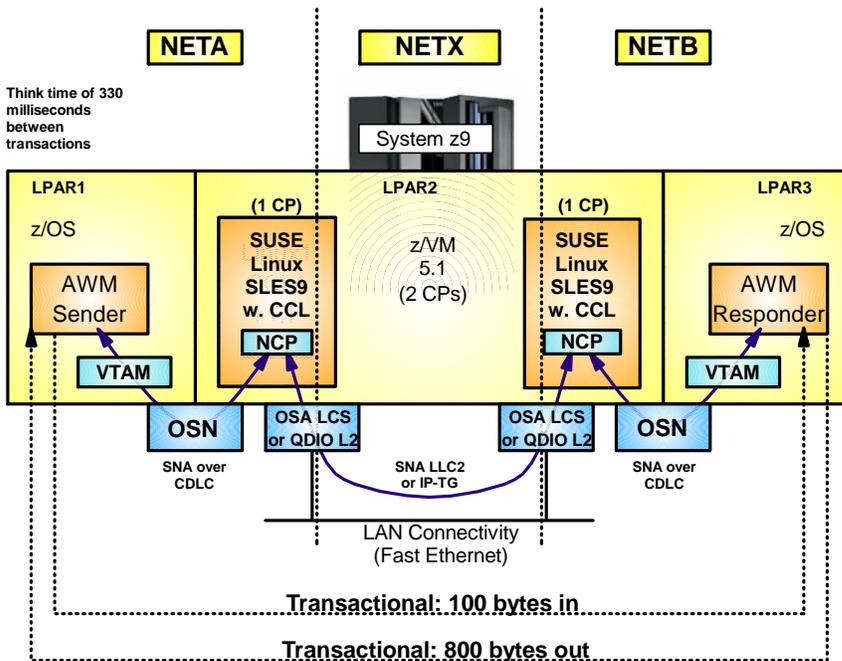
## CCL V1R2 – System z9 and zSeries CPU Capacity Planning



**Figure 2: IBM CCL SNI test environment using a shared LAN between VTAM and CCL**

This scenario is referred to as FE/FE – Fast Ethernet between VTAM and CCL – and Fast Ethernet between the two CCL NCPs.

The other test scenario for SNI workload used the OSA for NCP connectivity between VTAM and the CCL. In this scenario, z/OS and CCL need to reside on the same System z9 processor in order to share an OSA port configured in OSN mode.



**Figure 3: IBM CCL SNI test environment using OSN between VTAM and CCL**

## CCL V1R2 – System z9 and zSeries CPU Capacity Planning

This scenario is referred to as either CDLC/FE or CDLC/IPTG. CDLC/FE stands for CDLC over OSN between VTAM and CCL – and Fast Ethernet (SNA LLC2) between the two CCL NCPs. CDLC/IPTG stands for CDLC over OSN between VTAM and CCL – and IP Transmission Group (also over a Fast Ethernet in this case) between the two CCL NCPs.

In both CCL SNI test scenarios, the two Linux images running CCL were deployed as z/VM 5.1 guests, each with one CP. The two Linux images each had an OSA LCS port or a QDIO Layer-2 port.

For the test scenario using IP Transmission Group, each of the two Linux images had an OSA QDIO port connected to a 1 Gbps Ethernet instead of the usual OSA LCS ports connected to a Fast Ethernet.

For the test cases using LAN connectivity between z/OS and CCL, each of the two z/OS systems and each of the two Linux systems had separate OSA CHPIDs – for a total of four OSA CHPIDs.

### IBM 3745/46 boundary function test environment

For the boundary function tests with IBM 3745/46, a 16 Mbps token-ring LAN was used with the IBM 3745/46s each attached to that LAN using two TIC2 or two TIC3 adapters. As for the SNI test runs, two runs were performed with the BF workload: one using TIC2 and the other using TIC3 token-ring adapters.

Preliminary tests showed that using a single token-ring adapter for this type of workload limited the total throughput too much to establish a reasonable comparison with the CCL environment.

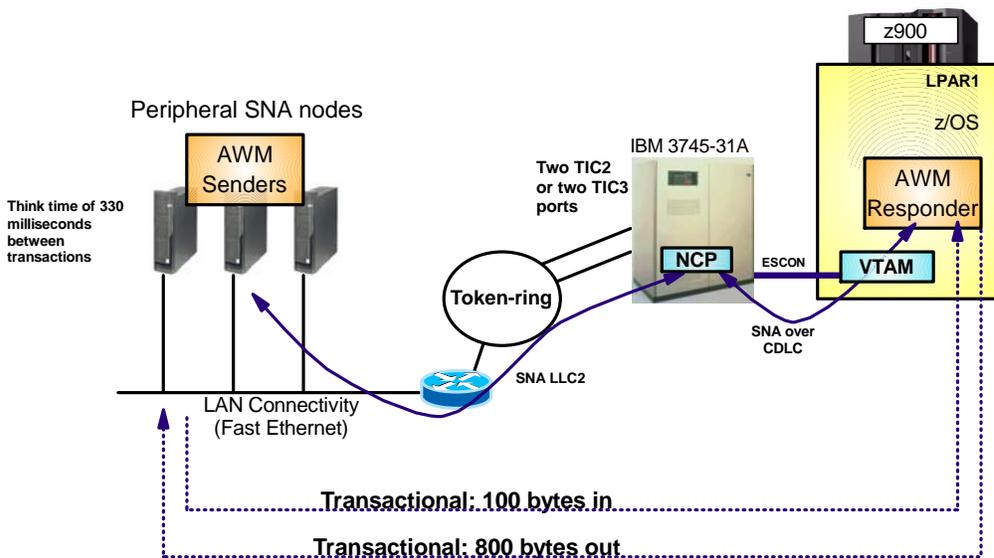
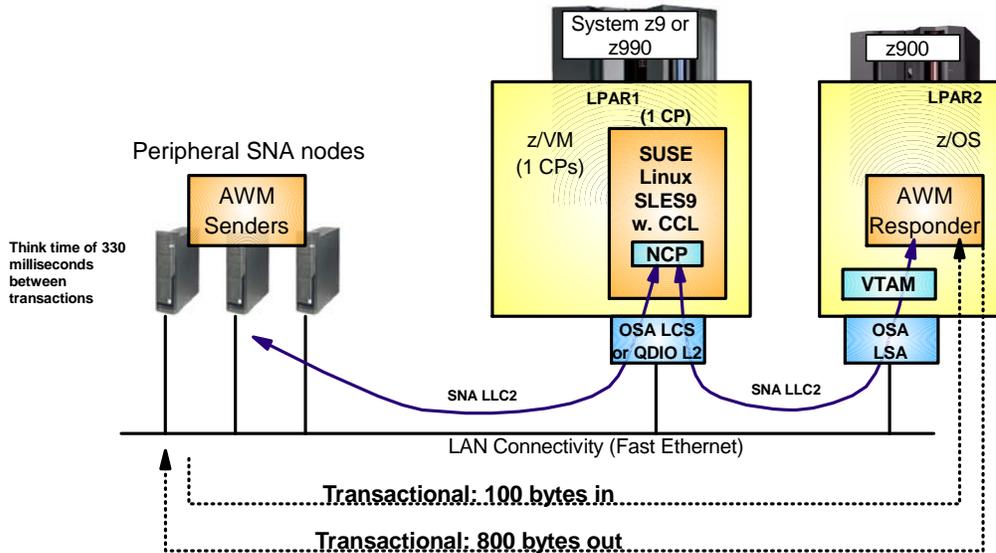


Figure 4: IBM 3745/46 boundary function test environment

### CCL boundary function test environment

For the boundary function tests with CCL, two test scenarios were deployed. The first scenario had z/OS connected to CCL over a Fast Ethernet (100 Mbps) LAN (z/OS using an OSA LSA port and CCL an OSA LCS port).

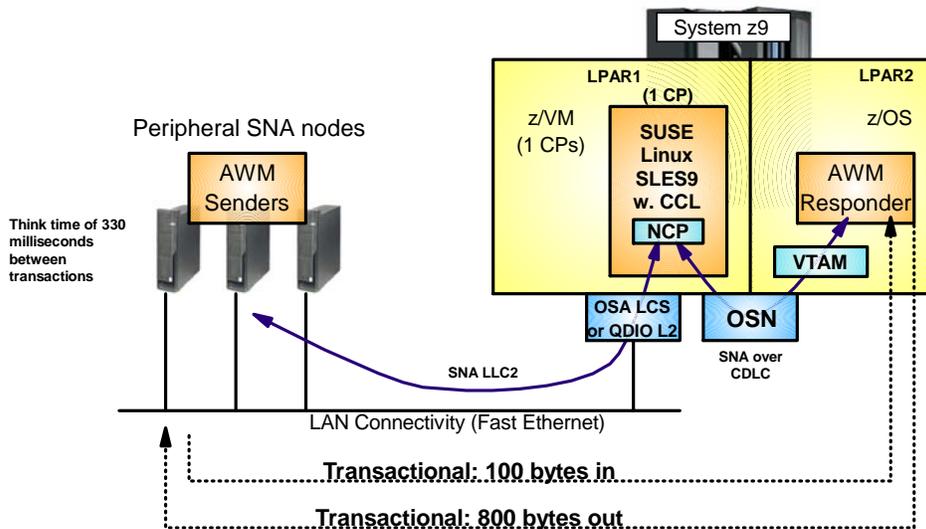
## CCL V1R2 – System z9 and zSeries CPU Capacity Planning



**Figure 5: IBM CCL boundary function test scenario using a shared LAN between VTAM and CCL**

This scenario is referred to as FE/FE – Fast Ethernet between VTAM and CCL – and Fast Ethernet between CCL and the peripheral nodes.

The other boundary function test scenario had VTAM and CCL connected via a shared OSA port operating in OSA for NCP mode. In this scenario, z/OS and CCL resided on the same System z9 processor in order to share an OSA port in OSN mode.



**Figure 6: IBM CCL boundary function test scenario using OSN between VTAM and CCL**

This scenario is referred to as CDLC/FE – CDLC over OSN between VTAM and CCL – and Fast Ethernet between CCL and the peripheral nodes.

The peripheral SNA nodes were connected to the CCL NCP through an OSA Ethernet port operating in OSA LCS mode or in QDIO Layer-2 mode. Each OSA port was assigned separate OSA CHPIDs – for a total of up to two OSA CHPIDs.

### ***CCL memory requirements***

The base CCL storage and memory requirements are:

- 55 MB DASD for Communication Controller for Linux and Java code
- 300 MB DASD for Linux kernel source
- 80 MB to 100 MB DASD per Communication Controller for Linux instance for traces, dumps, logs, and NCP load modules
- 20 MB RAM per Communication Controller for Linux instance

Our test environment was not storage constrained. Each z/VM guest had a virtual machine size of 512MB defined.

z/VM reported a storage working set between 240 MB and 310 MB for the CCL V1R2 SNI test runs, while a working set between 220 and 300 MB was reported for the CCL V1R2 boundary function test runs.

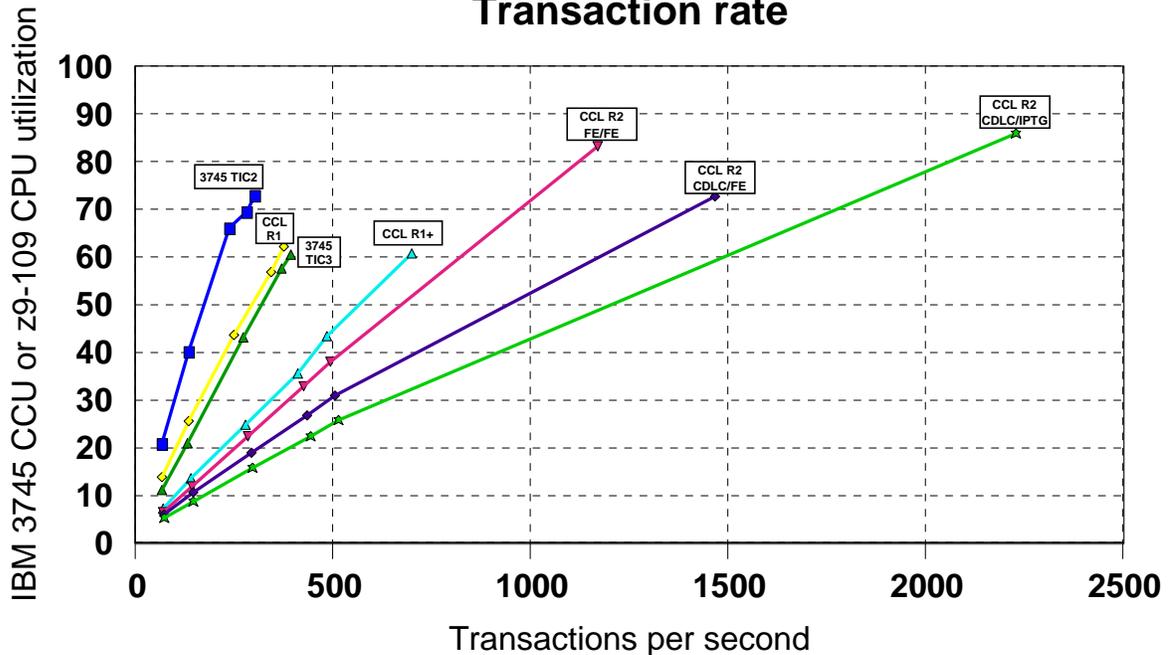
The test runs using OSN (CDLC between VTAM and CCL) and QDIO layer-2 in general used a higher working set – up to around 500 MB.

### SNI transactional workload - results

This section of the document describes the results of the SNI transactional workload test runs and provides basic capacity planning information for similar workloads.

The first comparison of the environments is based on raw throughput – how many transactions per second each of the environments can deliver.

## SNI Transactional Workload Measured throughput Transaction rate



The first observation to be made from this chart is that an IBM 3745/46 with TIC3 adapters provide higher throughput than an IBM 3745/46 with TIC2 adapters: for similar workloads, the IBM 3745 CCU needs to process more instructions when using TIC2 than when using TIC3. In our performance runs, an IBM 3745/46 with TIC2 adapters peak at around 300 transactions per second, while an IBM 3745/46 with TIC3 adapters peak at around 380 transactions per second.

The limiting factor for our IBM 3745-31A TIC3 test environment is not the CCU utilization, which at 380 transactions per second is just around 63% utilization. The limiting factor in the configuration that was used in this setup is the ESCON channel attachment to z/OS - which according to RMF measurements reached 85% utilization on one of the two z/OS systems, and the Token-Ring Processor (TRP) - which according to NTuneMON – also reached 85% utilization. With more ESCON channel connections and/or token-ring processor resources in the IBM 3745/46 configuration, we would most likely have been able to drive the transaction rate higher before the IBM 3745 CCU would have been exhausted. A rough extrapolation of the existing numbers suggest a rate of somewhere between 500 and 600 transactions per second as the limit.

The original CCL V1R1 code level performed better than an IBM 3745/46 with TIC2 adapters and very much similar to an IBM 3745/46 with TIC3 adapters.

## CCL V1R2 – System z9 and zSeries CPU Capacity Planning

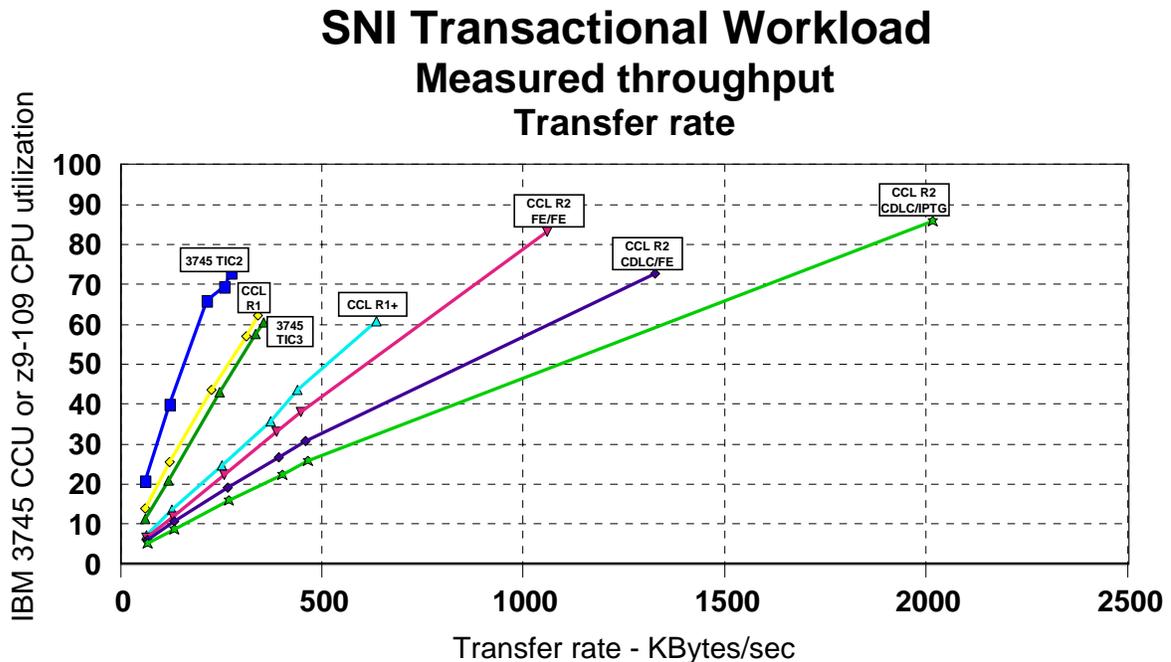
With the CCL V1R1+ code level from August 2005, the CCL environment generally outperforms the IBM 3745/46 environment.

With CCL V1R2 this becomes even more obvious, where CCL V1R2 provides transaction rates up to 2230 transactions per second – or nearly six times more than the measured transaction rate for an IBM 3745/46 environment using TIC3 adapters.

As the following charts will show, the limiting factor for CCL in our test runs is ultimately zSeries or System z9 CPU capacity.

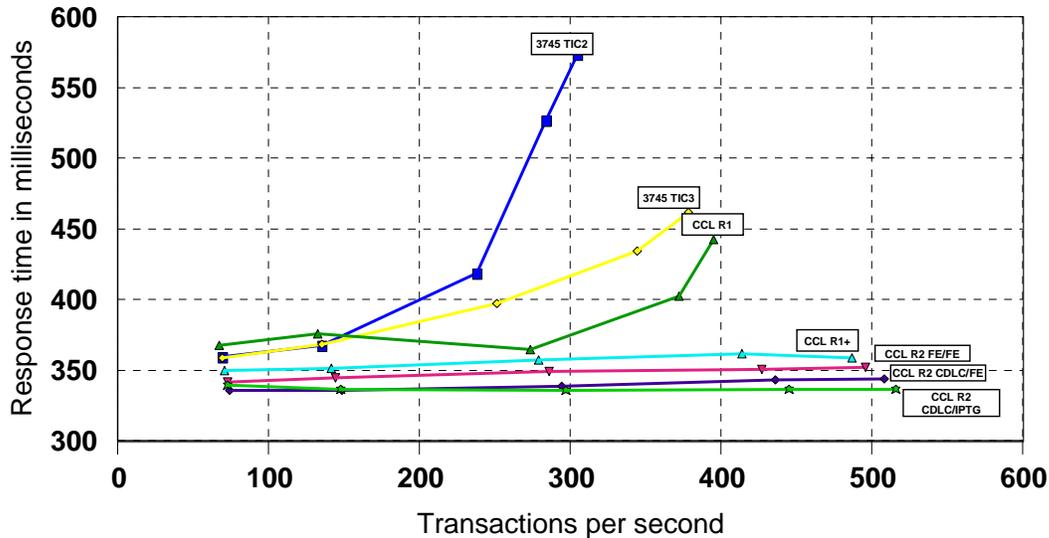
Another point to make about CCL V1R2 performance is that it does depend on the actual connectivity to CCL. OSN (CDLC) between VTAM and the CCL NCP, and IP Transmission Group (IPTG) between the two CCL NCPs provide the best performance.

Raw throughput can also be illustrated based on transfer rates – bytes per second.



The number of transactions per second cannot be viewed on its own; a comparison of how the response times vary also needs to be analyzed to obtain a more complete picture of the throughput. Since we used AWM as the test driver for our test runs, we do have response time data available as shown in the following chart.

## SNI Transactional Workload Measured response time



This chart illustrates how the environments compare from a response time perspective. As can be seen, the CCL environment delivers a better response time than the IBM 3745/46 environment – in particular compared to the IBM 3745/46 TIC2 environment.

This chart does not include the highest transaction rate measurements we captured during our test runs. To force our test driving system (AWM) to such high transaction rates, we had to perform those high transaction rate runs with a think time of zero, which obviously distorts the response time measurements. In the above chart, the response time includes the artificial 330 millisecond response time we in general used for all our test runs.

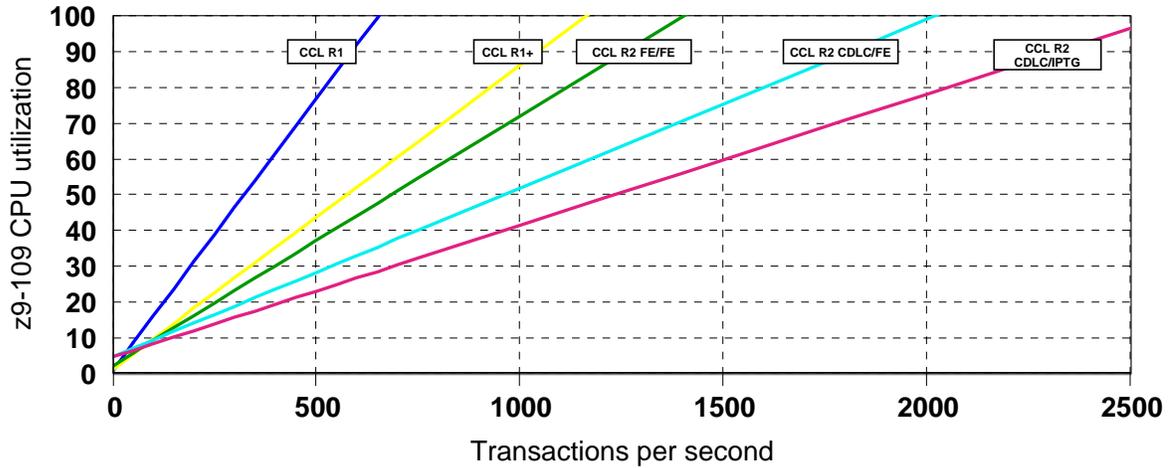
### Estimating SNI CCL CPU utilization

In the following charts, the estimate guidelines are based on transactions per second and existing IBM 3745-31A CCU utilization.

#### *Using transaction rates*

Given a transactional workload resembling the workload characteristics in our performance test, and given a known transaction rate, the following chart can be used to estimate the CCL CPU usage.

## SNI Transactional Workload z9-109 CPU utilization estimate Based on transaction rate



If the transaction rate is known, the CCL V1R2 CPU utilization on a System z9 can be estimated based on the following simple linear formulas:

**CCL V1R2 – FE between VTAM and CCL and FE between the two CCLs:**

$$\text{CCL CPU \%} = 2.25 + 0.0698 * \text{transaction rate}$$

**CCL V1R2 – CDLC between VTAM and CCL and FE between the two CCLs:**

$$\text{CCL CPU \%} = 4.62 + 0.0472 * \text{transaction rate}$$

**CCL V1R2 – CDLC between VTAM and CCL and IPTG between the two CCLs:**

$$\text{CCL CPU \%} = 4.56 + 0.0368 * \text{transaction rate}$$

### *Using IBM 3745 CCU utilization*

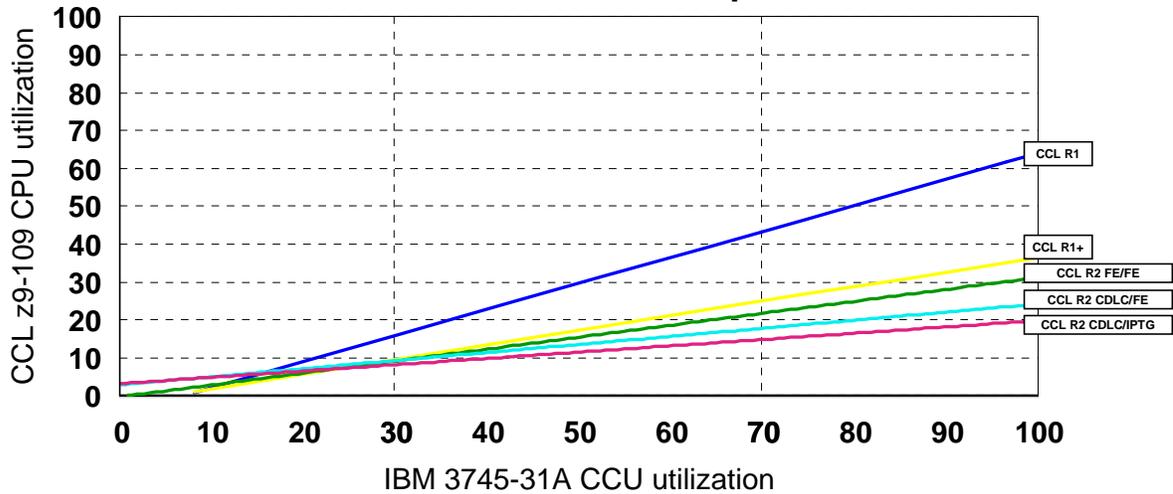
If the transaction rate is not known, or if the workload characteristics are unknown, then a rough estimate can be made based on the current IBM 3745 CCU utilization.

If you do use the current CCU utilization as a base for your CCL CPU utilization estimate, it is important to know whether the IBM 3745/46 environment uses TIC2 or TIC3 token-ring adapters, since the relationship between CCU utilization and System z9 CPU utilization is very different in those two cases.

The first chart shows the relationship between CCL System z9 CPU utilization and IBM 3745 CCU utilization in the case where the IBM 3745 NCP uses TIC2 adapters to connect to the LAN environment.

## SNI Transactional Workload

### IBM 3745-31A CCU utilization vs. z9-109 CPU utilization IBM 3745-31A with TIC2 adapters



The interval between 30% and 70% CCU utilization is the interval in which we expect most IBM 3745/46 configurations currently to operate.

For the TIC2 case, the formulas to estimate CCL V1R2 System z9 CPU utilization are:

**CCL V1R2 – FE between VTAM and CCL and FE between the two CCLs:**

$$\text{CCL CPU \%} = -0.315 + 0.319 * \text{IBM 3745 CCU utilization}$$

**CCL V1R2 – CDLC between VTAM and CCL and FE between the two CCLs:**

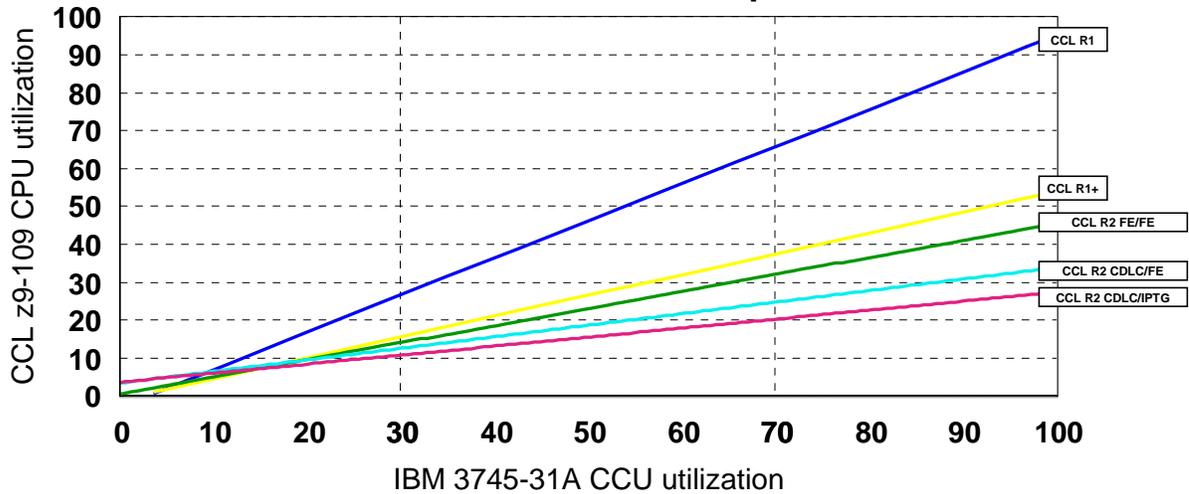
$$\text{CCL CPU \%} = 2.89 + 0.213 * \text{IBM 3745 CCU utilization}$$

**CCL V1R2 – CDLC between VTAM and CCL and IPTG between the two CCLs:**

$$\text{CCL CPU \%} = 3.2 + 0.166 * \text{IBM 3745 CCU utilization}$$

The second chart shows the relationship between CCL System z9 CPU utilization and IBM 3745 CCU utilization in the case where the IBM 3745 NCP uses TIC3 adapters to connect to the LAN environment.

## SNI Transactional Workload IBM 3745-31A CCU utilization vs. z9-109 CPU utilization IBM 3745-31A with TIC3 adapters



For the TIC3 case, the formulas to estimate CCL V1R2 System z9 CPU utilization are:

CCL V1R2 – FE between VTAM and CCL and FE between the two CCLs:

$$\text{CCL CPU \%} = 0.533 + 0.45 * \text{IBM 3745 CCU utilization}$$

CCL V1R2 – CDLC between VTAM and CCL and FE between the two CCLs:

$$\text{CCL CPU \%} = 3.46 + 0.304 * \text{IBM 3745 CCU utilization}$$

CCL V1R2 – CDLC between VTAM and CCL and IPTG between the two CCLs:

$$\text{CCL CPU \%} = 3.65 + 0.237 * \text{IBM 3745 CCU utilization}$$

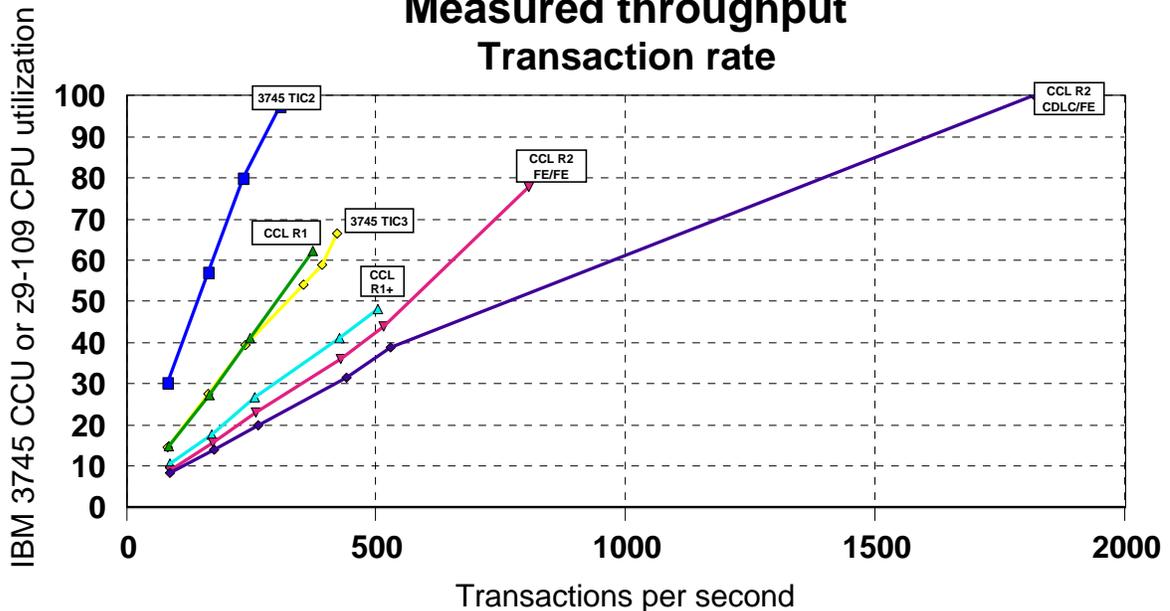
### Boundary function transactional workload - results

This section of the document describes the results of the BF transactional workload test runs and provides basic capacity planning information for similar workloads.

Please note that the measurement methodology we used for the BF workload has changed since the original CCL capacity planning reports were published. For this capacity planning report, all the CCL V1R1 and CCL V1R1+ measurements were repeated using the new measurement methodology to make the numbers comparable to the CCL V1R2 numbers, but there are in some cases minor discrepancies if you compare to the transaction rate numbers in the previous CCL capacity planning reports. We originally used the IBM Tele Processing Network Simulator (TPNS) product to drive the BF workload, but had to abandon that approach due to limitations in the driving systems. The new measurement methodology uses AWM as the driving system and uses message lengths and think times that are comparable to those we use for the SNI transactional workload.

The first comparison of the test environments is based on raw throughput – how many transactions per second each of the two environments can deliver.

## Boundary Function Transactional Workload Measured throughput Transaction rate

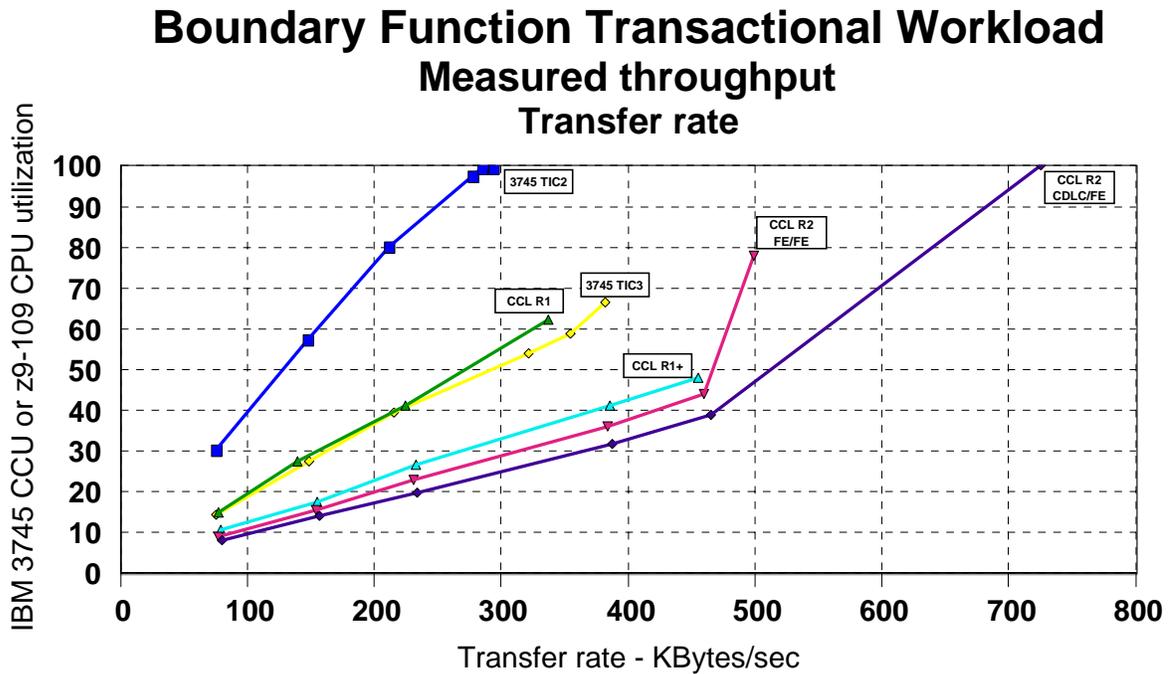


The IBM 3745-31A with TIC2 adapter environment was limited by CCU capacity.

Again, as was the case for the SNI workload runs, the IBM 3745-31A with TIC3 adapter environment was limited by ESCON channel and token-ring adapter capacity, and not by IBM 3745-31A CCU capacity.

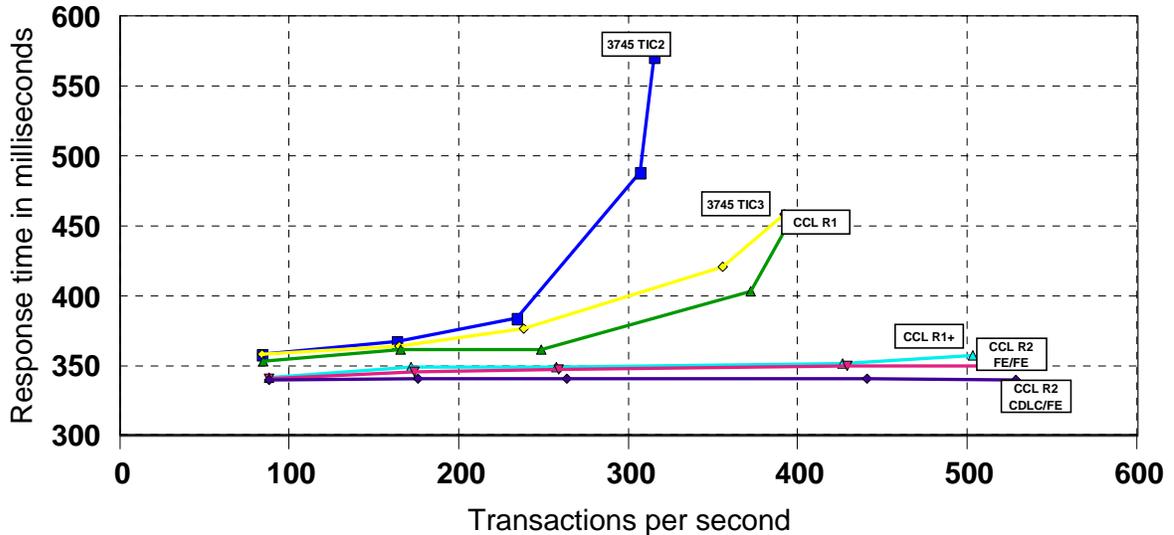
In order for the IBM 3745/46 environment to achieve these transactions rates, we had to configure two token-ring adapters. With a single token-ring adapter, the IBM 3745/46 environment wasn't able to achieve comparable transaction rates.

Raw throughput can also be illustrated based on transfer rates – bytes per second.



Again, as was the case for the SNI workload runs, our test driving system was initially not able to generate enough transactions per second to saturate the CCL CPU. Our workload characteristics were defined with an artificial 330 millisecond response time. In order to generate the highest transaction rates we had to include one run with a think time of zero, so the highest transaction rates are measured with a think time of zero, which doesn't distort the CPU utilization results, but would distort the response time measurements. For that reason, the highest transaction rate measurements are not included in the response time analysis chart, which follows:

## Boundary Function Transactional Workload Measured response times



The response time shown above includes an artificial think time between transactions of 330 milliseconds.

CCL V1R2 response times are flat at least up to a transaction rate between 500 and 600 transactions per second.

### Estimating boundary function CCL CPU utilization

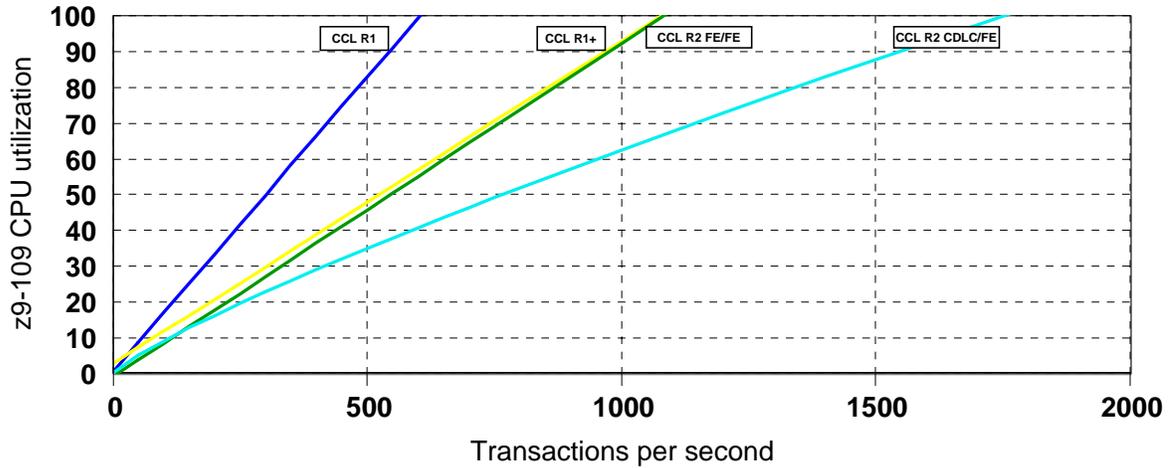
In the following charts, the estimate guidelines are based on transactions per second and existing IBM 3745-31A CCU utilization.

#### *Using transaction rates*

In this chart, the estimate guidelines are based on transactions per second.

If the workload is indeed transactional, assumed to have some resemblance to the workload characteristics used in our performance test runs, and the transaction rate is known – then the following chart and formulas can be used to estimate the CCL CPU usage.

## Boundary Function Transactional Workload z9-109 CPU utilization estimate Based on transaction rate



If the transaction rate is known, the CCL CPU utilization on a System z9 can be estimated based on the formulas:

**CCL V1R2 – FE between VTAM and CCL and FE between CCL and nodes:**

$$\text{CCL CPU \%} = -0.989 + 0.0935 * \text{transaction rate}$$

**CCL V1R2 – CDLC between VTAM and CCL and FE between CCL and nodes:**

$$\text{CCL CPU \%} = \text{transaction rate}^{0.839} * 0.19$$

Please note that the formula for estimating the CCL V1R2 System z9 CPU utilization based on transaction rate is a power formula and not a linear formula. Linear regressions on the measurement data did not provide as good a fit as a power formula did, so we chose to use the above power formula.

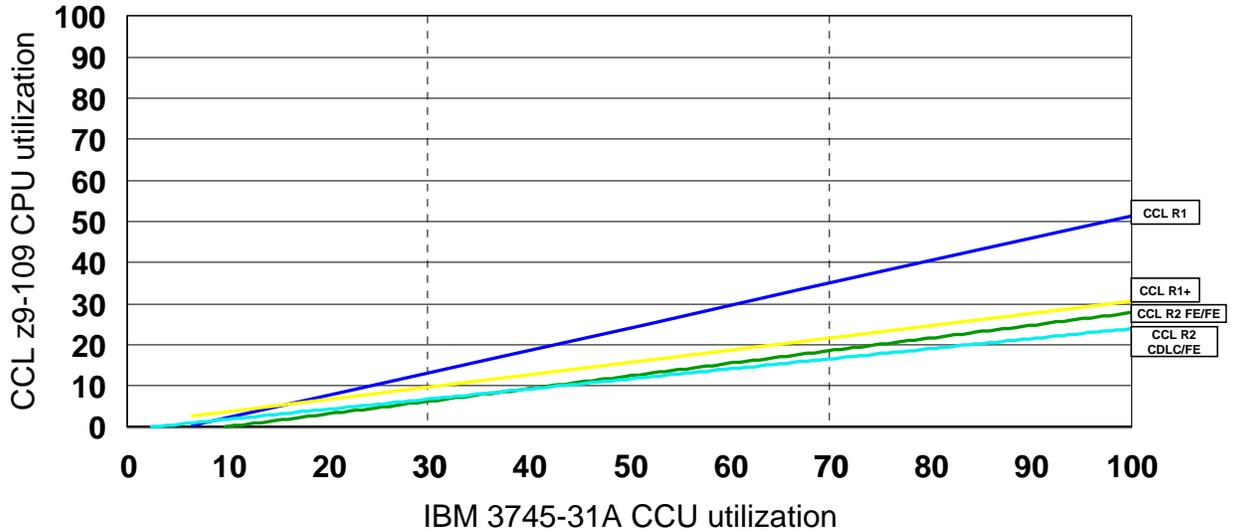
### *Using IBM 3745 CCU utilization*

If neither transaction rate nor transfer rate is known, or the workload characteristics are unknown, a rough estimate can be done based on the current IBM 3745 CCU utilization.

As for SNI workload, we also for boundary function workload need to distinguish between the cases where the original IBM 3745/46 environment uses TIC2 and TIC3 adapters.

The first chart shows the relationship between the IBM 3745-31A CCU utilization and the CCL V1R2 System z9 CPU utilization in the case where the IBM 3745/46 uses TIC2 adapters:

## Boundary Function Transactional Workload IBM 3745-31A CCU utilization vs. z9-109 CPU utilization IBM 3745-31A with TIC2 adapters



For the TIC2 case, the formulas to estimate System z9 CPU utilization for CCL V1R2 are:

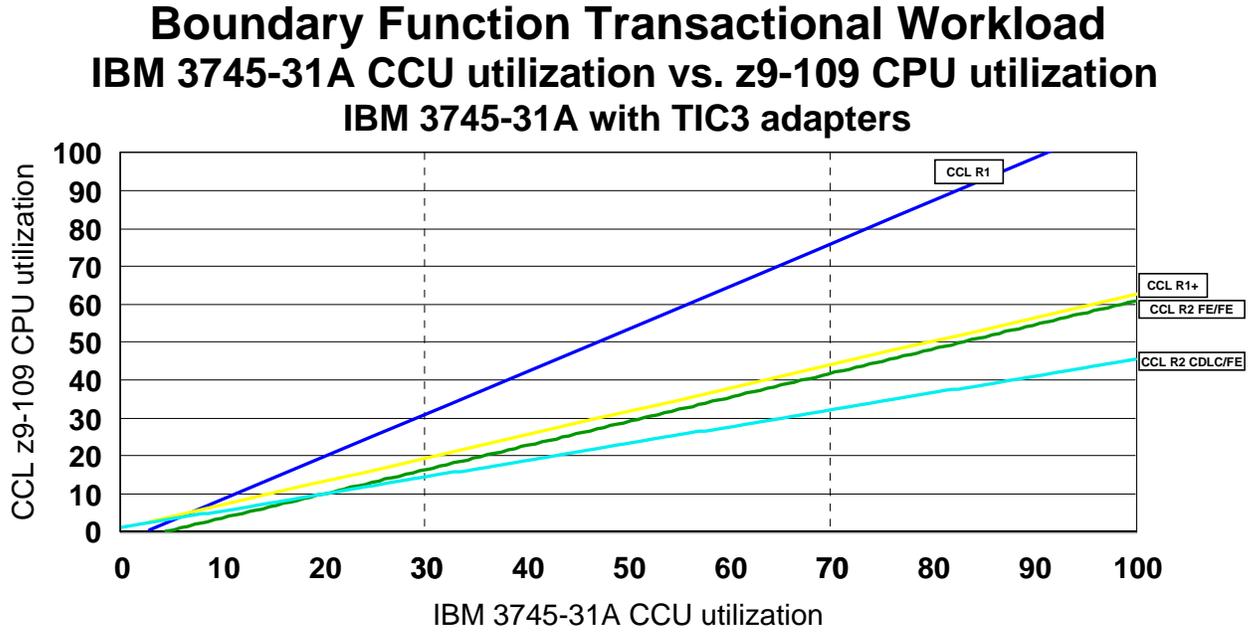
**CCL V1R2 - FE between VTAM and CCL and FE between CCL and nodes:**

$$\text{CCL CPU \%} = -3.004 + 0.308 * \text{IBM 3745 CCU utilization}$$

**CCL V1R2 - CDLC between VTAM and CCL and FE between CCL and nodes:**

$$\text{CCL CPU \%} = 0.451 + 0.226 * \text{IBM 3745 CCU utilization}$$

The next chart shows the relationship when the IBM 3745/46 uses TIC3 adapters:



For the TIC3 case, the formulas to estimate System z9 CPU utilization for CCL V1R2 are:

**CCL V1R2 - FE between VTAM and CCL and FE between CCL and nodes:**

$$\text{CCL CPU \%} = -2.856 + 0.638 * \text{IBM 3745 CCU utilization}$$

**CCL V1R2 - CDLC between VTAM and CCL and FE between CCL and nodes:**

$$\text{CCL CPU \%} = 0.561 + 0.469 * \text{IBM 3745 CCU utilization}$$

**Quick CCL V1R2 CPU utilization conversion tables**

If you know your current IBM 3745/46 CCU utilization and you know if it is equipped with TIC2 or TIC3 adapters, then use the following tables to get a rough estimate of the CCL V1R2 System z9 CPU utilization.

Locate the table that matches your type of workload (SNI or BF) and your type of TIC adapters (TIC2 or TIC3), locate your current CCU utilization row, and pick the System z9 CPU utilization number in the column that matches your planned CCL configuration.

Example: if you are using an IBM 3745/46 with TIC2 adapters for transactional SNI workload and your current CCU utilization is 50%, then your estimated CCL V1R2 CDLC/IPTG System z9 CPU utilization will be 11.51%.

**SNI from IBM 3745-31A with TIC2 to z9-109**

<b>IBM 3745 TIC2 CCU Utilization</b>	<b>CCL V1R2 FE/FE</b>	<b>CCL V1R2 CDLC/FE</b>	<b>CCL V1R2 CDLC/IPTG</b>
10	2.84	5.01	4.87
20	5.98	7.14	6.53
30	9.13	9.27	8.19
40	12.28	11.40	9.85
50	15.43	13.53	11.51
60	18.58	15.66	13.17
70	21.73	17.79	14.83
80	24.88	19.92	16.49
90	28.03	22.05	18.15
100	31.18	24.18	19.82

**SNI from IBM 3745-31A with TIC3 to z9-109**

IBM 3745 TIC3 CCU Utilization	CCL V1R2 FE/FE	CCL V1R2 CDLC/FE	CCL V1R2 CDLC/IPTG
10	5.04	6.50	6.03
20	9.54	9.55	8.40
30	14.04	12.59	10.78
40	18.54	15.63	13.15
50	23.04	18.68	15.52
60	27.54	21.72	17.90
70	32.04	24.76	20.27
80	36.54	27.81	22.65
90	41.04	30.85	25.02
100	45.54	33.89	27.39

**BF from IBM 3745-31A with TIC2 to z9-109**

IBM 3745 TIC2 CCU Utilization	CCL V1R2 FE/FE	CCL V1R2 CDLC/FE
10	0.08	2.03
20	3.16	4.42
30	6.24	6.81
40	9.32	9.20
50	12.40	11.59
60	15.48	13.98
70	18.56	16.37
80	21.64	18.76
90	24.72	21.15
100	27.80	23.54

**BF from IBM 3745-31A with TIC3 to z9-109**

IBM 3745 TIC3 CCU Utilization	CCL V1R2 FE/FE	CCL V1R2 CDLC/FE
10	3.52	5.61
20	9.90	10.00
30	16.27	14.39
40	22.65	18.78
50	29.03	23.17
60	35.40	27.56
70	41.78	31.95
80	48.16	36.34
90	54.53	40.73
100	60.91	45.12

**Notes**

All the CCL CPU utilization numbers in the above measurements were based on the z/VM reported CPU utilization that includes z/VM processing overhead.

All measurements that were used in this document were obtained running Linux as guests under z/VM. Additional measurements in our test environment show that if you deploy Linux in an LPAR without z/VM, you should expect performance numbers that are comparable to or slightly better than those found in a z/VM environment.

In the SNI test environment, one z/OS system and one NCP constituted the client side, while the other NCP and z/OS system constituted a server side. There were slight differences in the CPU utilization between the client side NCP and the server side NCP, but they were minor (within 2 percentage points for CCL and 4 for the IBM 3745-31A). The numbers used in the above charts are the average of the client and server side NCP utilization numbers.

Each CCU in an IBM 3745-31A/61A is equivalent to the CCU in an IBM 3745-31A. If you are using another IBM 3745 model you can use the following relative throughput ratios to adjust your CCU utilization to the IBM 3745-31A utilization data used in this document.

IBM 3745 CCU	Relative Throughput Ratio
IBM 3745-17A	<b>1.00</b>
IBM 3745-21A/41A	<b>1.66</b>
IBM 3745-31A/61A	<b>2.50</b>

Do keep in mind that your workload may vary significantly from the workload described in this document. You should be careful using the above guidelines to estimate your CCL CPU utilization.

Please also note that at very low or very high IBM 3745 CCU utilization values, the estimated CCL CPU utilization tends to be less accurate than if your current IBM 3745 CCU utilization lies within the “normal” range of 30% to 70% CCU utilization.

## CCL V1R2 – System z9 and zSeries CPU Capacity Planning

If you are not using a System z9 CPU, please use the usual Large Systems Performance Ratio (LSPR) data to adjust the CPU capacity to the level of zSeries hardware you are planning to use. Please see the following site for LSPR information:

<http://www-03.ibm.com/servers/eserver/zseries/lSpr/>

In our case, we used a ratio 1:1.35 between a z990 and a System z9 CPU – the System z9 CPU being 35% faster than a z990 CPU.

### **Samples**

We'll go through two simple examples that illustrate how to use the above capacity planning information.

#### **IBM 3745 17A SNI workload to CCL V1R2 on a z990**

If you today use an IBM 3745 17A for transactional SNI workload, your IBM 3745 17A CCU is 70% utilized, your LAN connectivity is through TIC2 adapters, and you want to migrate that workload to CCL V1R2 running on a z990 using Fast Ethernet connectivity between VTAM and the CCL NCP and Fast Ethernet for downstream connectivity, then go through the following steps:

1. Start by converting the IBM 3745 17A utilization to IBM 3745 31A CCU utilization:  $70\% / 2.50 = 28.0\%$ .
2. Calculate your CCL V1R2 System z9 utilization by using the formula for SNI workload from TIC2 to FE/FE CCL connectivity

$$(\text{CCL CPU \%} = -0.315 + 0.319 * \text{IBM 3745 CCU utilization})$$

In our example:  $-0.315 + 0.319 * 28.0\% = 8.9\%$

3. Final step is to convert that System z9 CPU utilization to z990 CPU utilization:  $8.9\% / (1 - 0.35) = 13.7\%$

#### **IBM 3745 21A boundary function workload to CCL V1R2 on a System z9**

If you today use an IBM 3745 21A CCU for boundary function workload, your IBM 3745 21A CCU is 65% utilized, your LAN connectivity is through TIC3 adapters, and you want to migrate that workload to CCL V1R2 running on a System z9 using OSN connectivity between VTAM and the CCL NCP and Fast Ethernet for downstream connectivity, then go through the following steps:

1. Start by converting the IBM 3745 21A utilization to IBM 3745 31A CCU utilization:  $65\% / 2.5 * 1.66 = 43.2\%$
2. Calculate your CCL V1R2 System z9 utilization by using the formula for boundary function workload from TIC3 to CDLC/FE CCL connectivity

$$(\text{CCL CPU \%} = 0.561 + 0.469 * \text{IBM 3745 CCU utilization})$$

In our example:  $0.561 + 0.469 * 43.2\% = 20.8\%$