



IBM Software Group
Enterprise Networking and Transformation Solutions (ENTS)

IBM Communication Controller for Linux (CCL) on System z

A Technical Introduction

Alfred B Christensen, Raleigh,
North Carolina, USA
alfredch@us.ibm.com

Trademarks and notices

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both:

- | | | | |
|-------------------------------------|----------------------------|-------------------------|------------------|
| ▶ Advanced Peer-to-Peer Networking® | ▶ GDDM® | ▶ OMEGAMON® | ▶ System i5 |
| ▶ AIX® | ▶ HiperSockets | ▶ Open Power | ▶ System p5 |
| ▶ alphaWorks® | ▶ HPR Channel Connectivity | ▶ OpenPower | ▶ System x |
| ▶ AnyNet® | ▶ HyperSwap | ▶ Operating System/2® | ▶ System z |
| ▶ AS/400® | ▶ i5/OS (logo) | ▶ Operating System/400® | ▶ System z9 |
| ▶ BladeCenter® | ▶ i5/OS® | ▶ OS/2® | ▶ Tivoli (logo)® |
| ▶ Candle® | ▶ IBM (logo)® | ▶ OS/390® | ▶ Tivoli® |
| ▶ CICS® | ▶ IBM® | ▶ OS/400® | ▶ VTAM® |
| ▶ DB2 Connect | ▶ IMS | ▶ Parallel Sysplex® | ▶ WebSphere® |
| ▶ DB2® | ▶ IP PrintWay | ▶ PR/SM | ▶ xSeries® |
| ▶ DRDA® | ▶ IPDS | ▶ pSeries® | ▶ z9 |
| ▶ e-business on demand® | ▶ iSeries | ▶ RACF® | ▶ zSeries® |
| ▶ e-business (logo) | ▶ LANDP® | ▶ Rational Suite® | ▶ z/Architecture |
| ▶ e business (logo)® | ▶ Language Environment® | ▶ Rational® | ▶ z/OS® |
| ▶ ESCON® | ▶ MQSeries® | ▶ Redbooks | ▶ z/VM® |
| ▶ FICON® | ▶ MVS | ▶ Redbooks (logo) | ▶ z/VSE |
| | ▶ NetView® | ▶ Sysplex Timer® | |

- Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- Intel, Intel Inside (logos), MMX and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
- Red Hat is a trademark of Red Hat, Inc.
- SUSE® LINUX Professional 9.2 from Novell®
- Other company, product, or service names may be trademarks or service marks of others.
- This information is for planning purposes only. The information herein is subject to change before the products described become generally available.
- All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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

Refer to www.ibm.com/legal/us for further legal information.

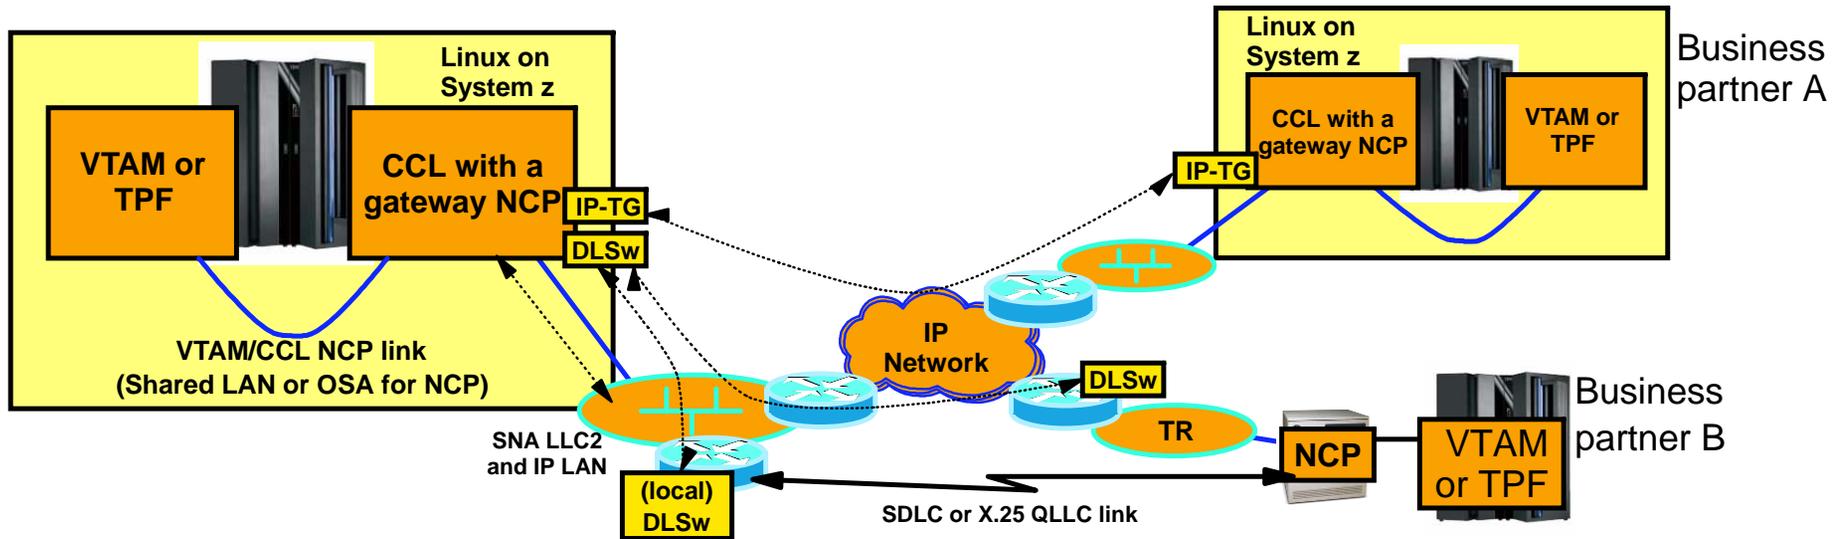
Agenda



1. CCL Objectives and Overall Technical Structure
2. CCL Connectivity Options
3. Performance and Capacity
4. Summary
5. System Requirements and Reference Information

CCL Objectives and Overall Technical Structure

Primary objective of CCL: preserve SNI connectivity to business partners (SNI/INN traffic)



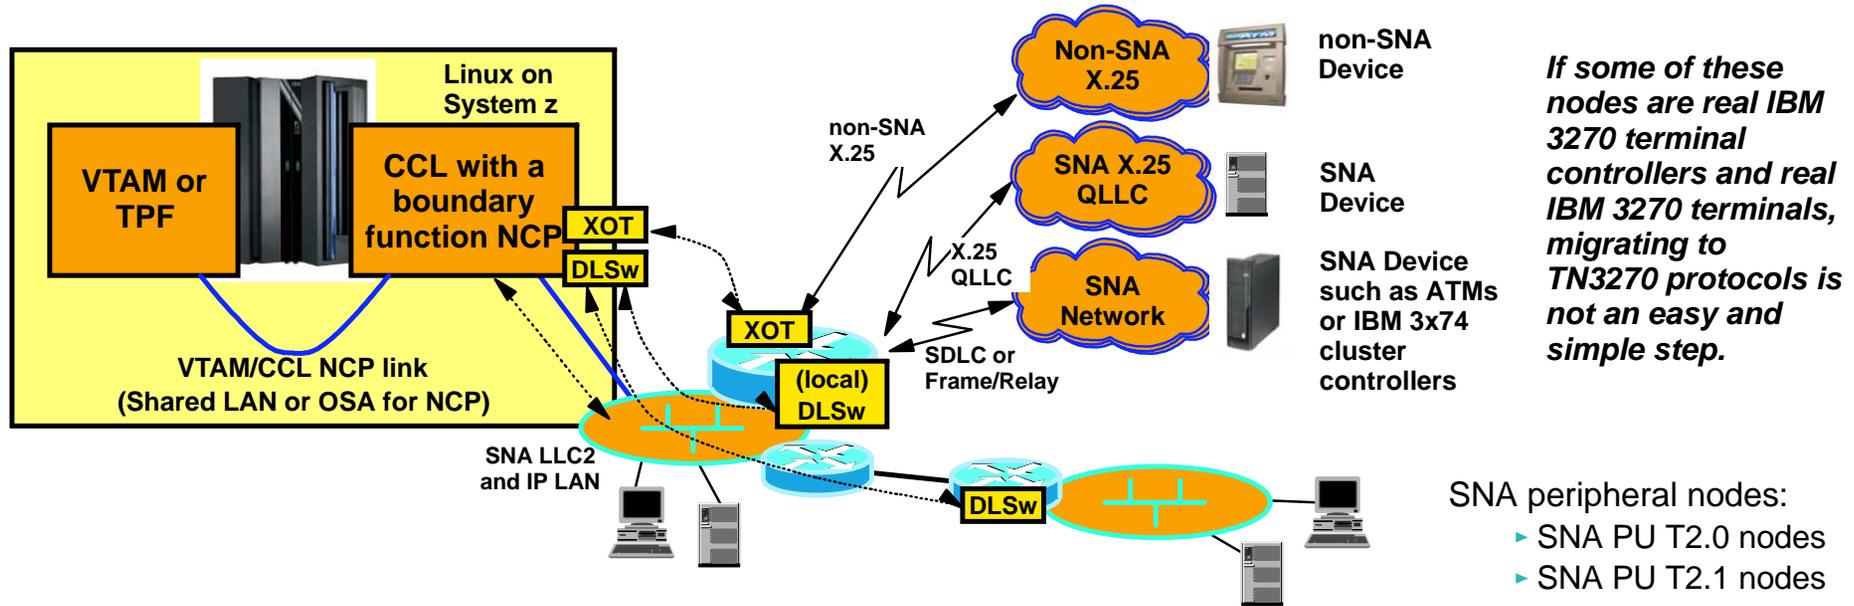
➤ Preserve existing SNA subarea capabilities and topology for business partner connectivity:

- ▶ For business partner SNA connectivity where your partner or you have decided to continue such communication based on SNA subarea networking protocols (SNI).
 - Business partners may continue to use IBM 3745/46 technology or also move to a CCL implementation.
- ▶ Has no impact on existing SNI topology.
- ▶ Has minimal impact on existing SNA network management procedures and disciplines.

➤ Simplify the networking infrastructure by Integrating SNA and IP traffic over common IP-based network:

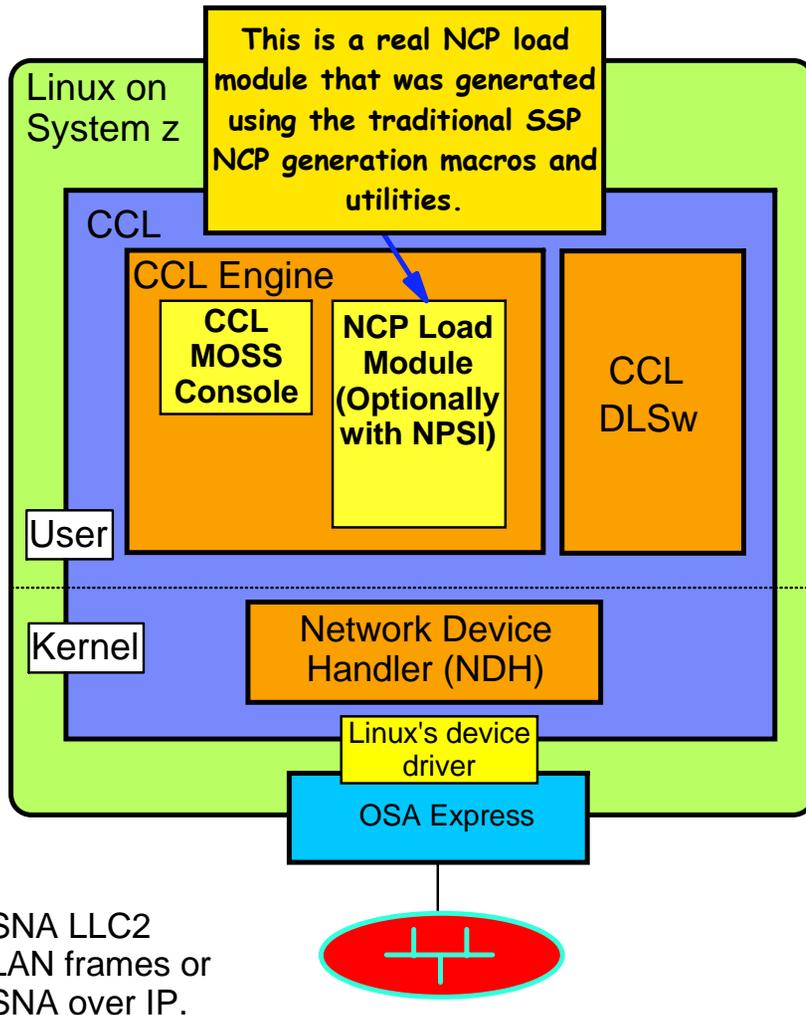
- ▶ CCL NCP to CCL NCP connectivity (INN or SNI).
 - SNI or INN over a TCP connection - IP Transmission Group (IP-TG)
- ▶ CCL NCP to IBM 3745 NCP connectivity (INN or SNI).
 - CCL imbedded DLSw to remote DLSw node that switches to TR to which partner IBM 3745/46 is connected
 - CCL LAN SNA LLC2 to local aggregation layer router in which SDLC or X.25 QLLC link to partner IBM 3745/46 is terminated

Secondary objective of CCL: preserve selected NCP boundary functions (BNN traffic)



- **Preserve existing SNA subarea capabilities and topology for peripheral node connectivity:**
 - ▶ NCP boundary function support includes standard availability functions such as SSCP takeover, support for duplicate MAC addressing, NRF, and XRF.
 - ▶ SNA network management tools and functions such as NetView, NPA-LU, and NtuneMON are also supported.
- **SNA serial line termination is supported via a network aggregation layer router:**
 - ▶ LLC bridging/switching of the SNA frames between the serial line interfaces and the local LAN
 - ▶ Serial line SNA connectivity for SDLC, Frame Relay, and SNA X.25 QLLC links are supported
- **Simplify the networking infrastructure by Integrating SNA and IP traffic over common IP-based network:**
 - ▶ Remote SNA network segments can be connected via DLSw over an IP network to the imbedded DLSw component in CCL
 - This includes both LAN segments, and local and remote aggregation layer routers
 - ▶ Non-SNA X.25 links are supported via X.25 Over TCP (XOT) for connectivity to NPSI

CCL overall structure and components



- **CCL** supports an **NCP** performing Boundary Functions, INN, and SNI link connectivity, as well as NPSI .
- **CCL** consists of both user-space and kernel-space functions:
 - ▶ **CCL engine** emulates an IBM 3745-31A with 16 MB memory supporting an NCP load module and a MOSS console interface.
 - ▶ The **MOSS console** is accessed through a standard Web browser.
 - ▶ **Network Device Handler (NDH)** is a kernel extension that acts as the interface between a real network interface (such as an OSA port) and the CCL adapter emulation support
 - ▶ **CCL DLSw** is a separate user-space application that communicates with CCL NCP through NDH, and with other DLSw peers through the Linux TCP sockets layer.

Note: You will continue to use ACF/SSP to generate, load, and dump an NCP load module.

CCL and the MOSS console interface

➤ The CCL MOSS console functions are accessed via a Web browser.

IBM.
Communication Controller for Linux

CCL Name: CCL44
NCP Name: CCL44
Machine Time: 01/27/2006 03:06:19 PM

Status	X71	X72	LAR	IAR	Level	C-Latch	Z-Latch
Running	000000	000000	1605CA	1605CC	6	0	0

Disk IPL Information

CP Running: CCL44
Auto Dump/Load: Yes
Active Load Module: CCL44

Disk Contents:

Type:	Name:	Save:	Gen:	IPL:	IPL Alert:
Load Module	CCL44	12/13/2005 04:27:41 PM	12/13/2005 03:04:40 PM	None	None

Dump

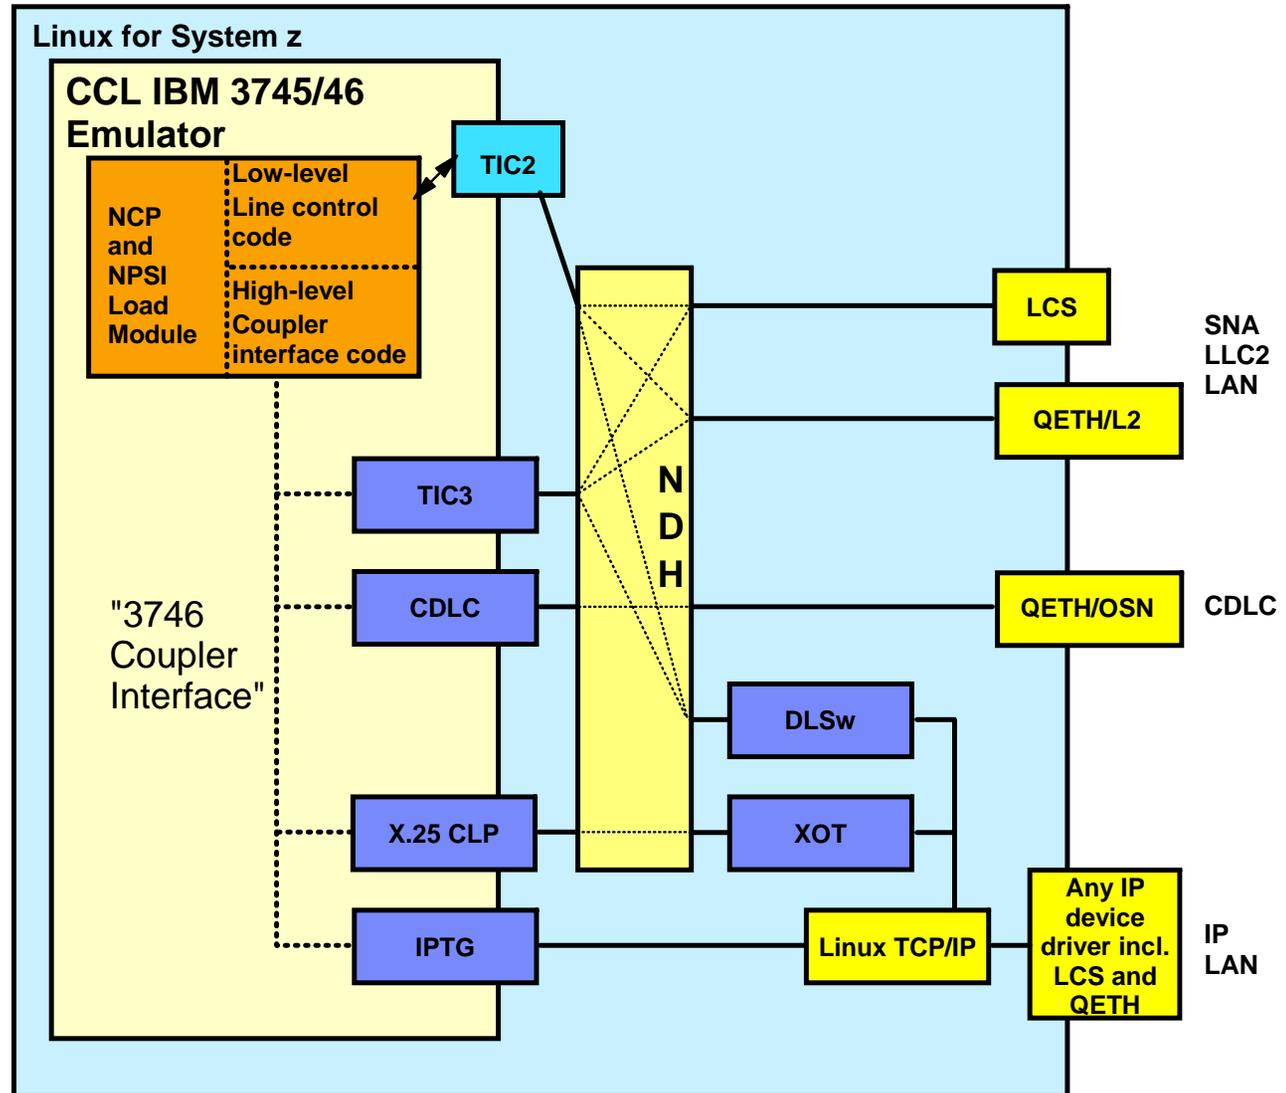
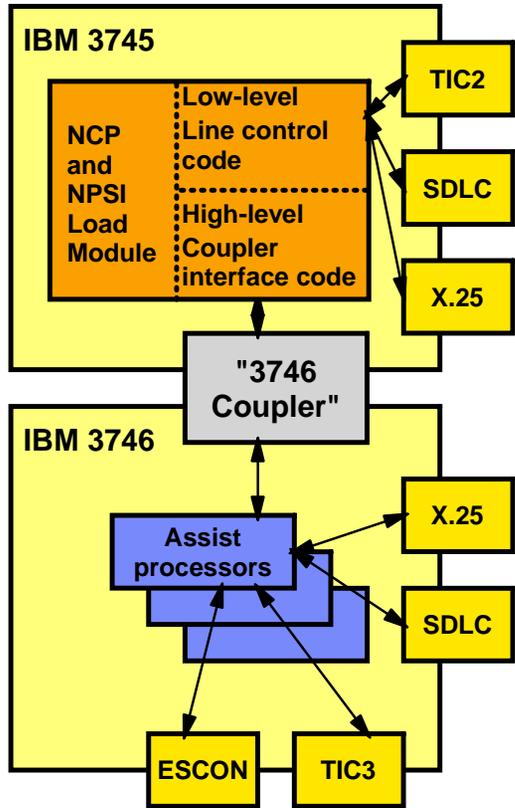
Purge Dump Change Dump/Load Change Active Load Module Rename Load Modules

© 2004, 2005 International Business Machines Corporation. All rights reserved.

MOSS console access is password protected!

CCL Connectivity Options

CCL NCP network interface architecture



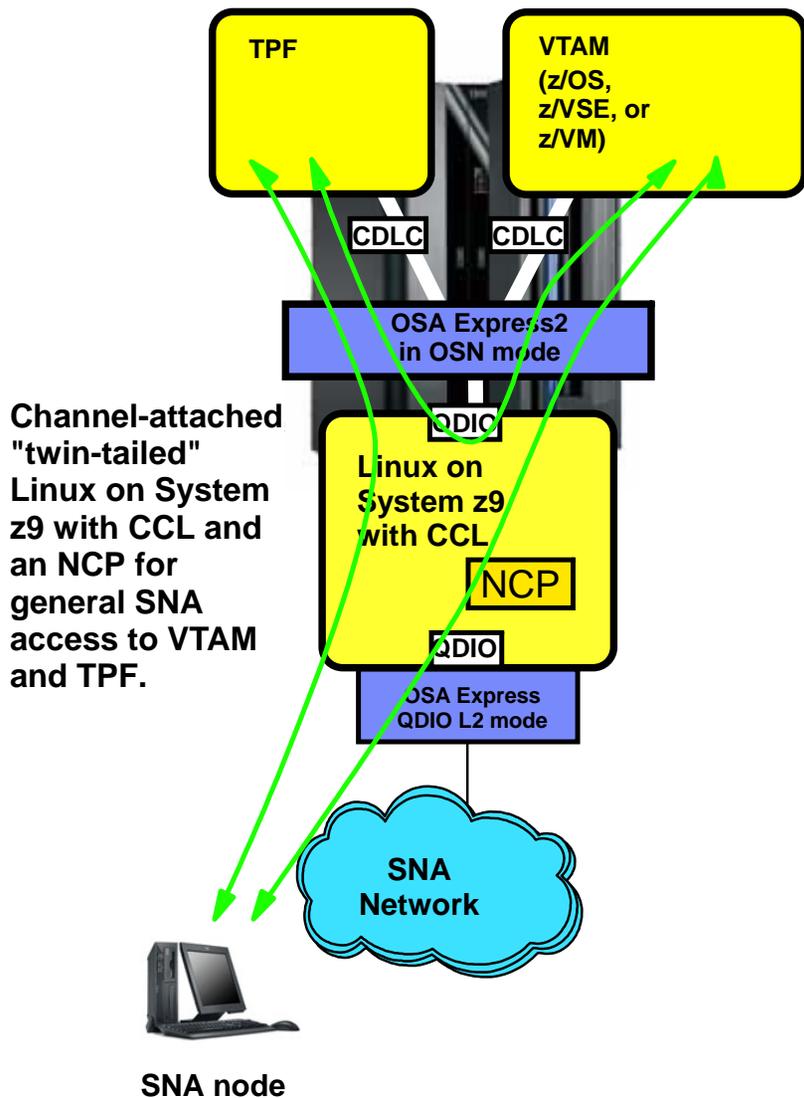
➤ Using network interfaces that are accessed through the IBM 3746 coupler interface, offloads the low-level line-specific control functions from the NCP to one of the assist processors in the IBM 3746 frame.

➤ For an NCP running in the CCL emulator - that means improved performance:

1. Less instructions to process by the emulator
2. Improved multi-processing capabilities - handing work from the emulator process to other processes and threads in Linux

- CCL V1R1 supported only the TIC2 (NTRI) technology over OSA LCS
- CCL V1R2 added support for CDLC, IPTG, X.25/XOT, and TIC2 over OSA QETH (QDIO in Layer 2 mode)
- CCL V1.2.1 added support for DLSw and TIC3 over LCS or QETH (QDIO in Layer 2 mode)

OSA for NCP (OSN) connectivity on System z9 - CDLC



- **OSA Express2 on System z9 implements an OSA CHPID type - known as OSA for NCP (OSN).**
 - OSA-Express2 1000Base-T and 1 Gigabit Ethernet features on z9-109
- **TPF and VTAM see the OSA Express OSN port as a channel-attached IBM 3745 to which they communicate using the usual CDLC channel protocol.**
 - OSA CHPID defined as OSN
 - TPF and VTAM device number defined as IBM 3745
 - Linux device numbers (three in a set) defined as OSN (accessed through QDIO from Linux)
 - Same OSN port can be shared among more SNA host systems and more CCL NCPs
- **The OSA microcode relays the CDLC data over a QDIO interface to CCL, which presents it to the NCP as though it had arrived over an IBM 3746 channel interface.**
 - No cable needs to be attached to an OSA port that is defined in OSN mode
- **The fact that the NCP runs in CCL instead of an IBM 3745/46 is transparent to TPF and VTAM.**
 - Existing configuration definitions are used unchanged
 - Existing activation and management flows continue to work as before
- **The normal Load/Dump functions over a channel are supported**
 - No need to FTP an NCP load module to the Linux file system
- **TPF or VTAM must reside on the same System z9 CEC as where CCL resides**
 - This is a same-CEC connectivity technology

CCL SNA LLC2 LAN connectivity overview

➤ NCP definitions for two types of LAN adapters:

- ▶ **TIC2** - NCP Token-ring Interface (NTRI) support
 - NCP instructions used for LLC2 logic
- ▶ **TIC3** - IBM 3746 Token-ring Processor (TRP)
 - Native System z instructions used for LLC2 logic
 - Also referred to as "native LAN" support
 - TIC3 performs better than TIC2 - uses much less CPU

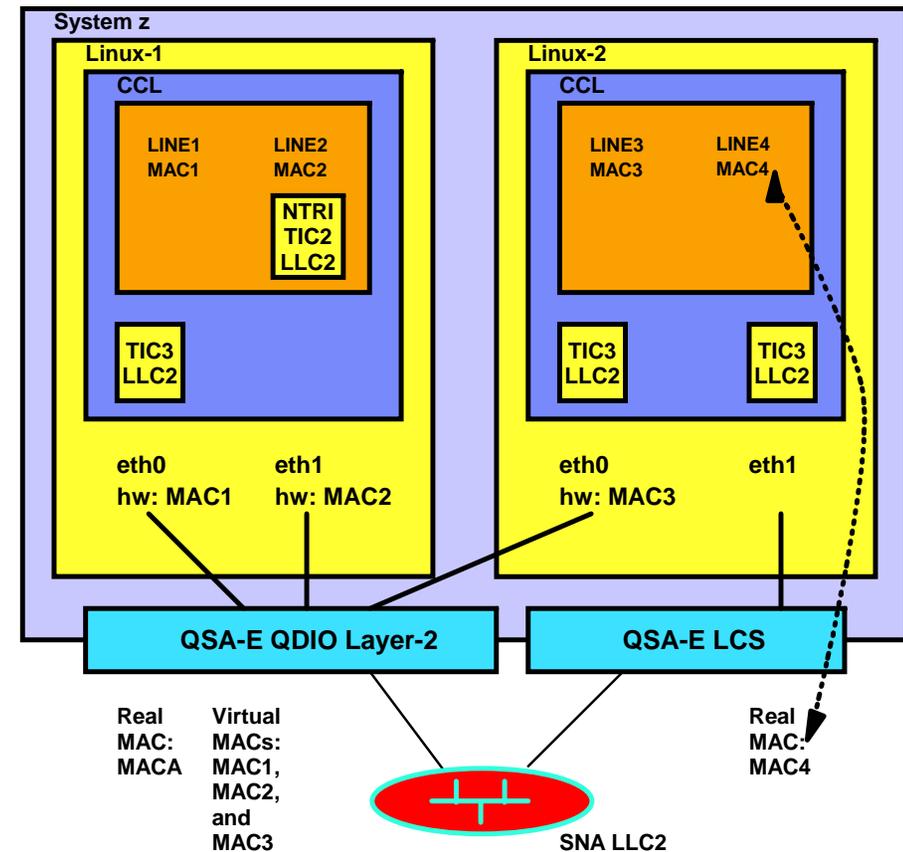
➤ CCL uses two different LAN device drivers for SNA LLC2:

▶ Lan Channel Station (LCS)

- Copper cabling
- NCP physical LINE local MAC address specification must match OSA port's configured real MAC address (OSA/SF)
- One NCP physical line per OSA LCS port
- Limited sharing via configured SAP numbers

▶ Queued Direct IO (QDIO) operating in layer 2 mode

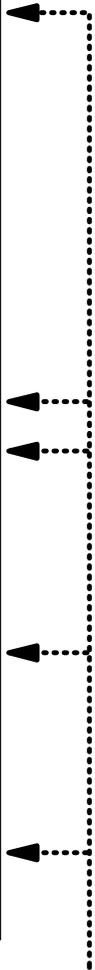
- Copper and fiber cabling
- NCP physical LINE local MAC address specification must match Linux interface hardware address (virtual MAC), but not OSA port's real MAC address
 - Up to 2048 virtual MAC addresses per OSA port
- Requires a Linux 2.6 kernel
- Works for Linux LPARs or z/VM guests
- If used by Linux as z/VM guests it can be used in combination with z/VM's virtual switch



Layer-2 mode is supported by Fast Ethernet, 1000BASE-T Ethernet, Gigabit Ethernet, and 10 Gigabit Ethernet features on OSA-Express and OSA-Express2 on z890, z990, and System z9

OSA-Express connectivity overview

Feature	Feature Name	Ports	z800 z900	z900 z990	z9-109	CHPIDs	Connectors
5201	OSA-2 Token Ring	2	X	N / A	N / A	OSA	Copper, RJ-45
5202	OSA-2 FDDI	1	X	N / A	N / A	OSA	Fiber, SC Duplex
2362	OSA-E 155 ATM SM	2	X	RPQ	N / A	OSD, OSE	Fiber, SC Duplex
2363	OSA-E 155 ATM MM	2	X	RPQ	N / A	OSD, OSE	Fiber, SC Duplex
2364	OSA-E GbE LX	2	X	C	C	OSD L2/L3**	Fiber, SC Duplex
2365	OSA-E GbE SX	2	X	C	C	OSD L2/L3**	Fiber, SC Duplex
2366	OSA-E Fast Ethernet	2	X	C	C	OSD L2/L3**, OSE	Copper, RJ-45
2367	OSA-E Token Ring	2	X	X	N / A	OSD, OSE	Copper, RJ-45
1364	OSA-E GbE LX	2	09/04	06/03	C	OSD L2/L3**	Fiber, LC Duplex
1365	OSA-E GbE SX	2	09/04	06/03	C	OSD L2/L3**	Fiber, LC Duplex
1366	OSA-E 1000BASE-T Ethernet	2	N / A	06/03	C	OSC, OSD L2/L3**, OSE	Copper, RJ-45
3364	OSA-E2 GbE LX	2	N / A	01/05	X	OSD L2/L3**, OSN *	Fiber, LC Duplex
3365	OSA-E2 GbE SX	2	N / A	01/05	X	OSD L2/L3**, OSN *	Fiber, LC Duplex
3366	OSA-E2 1000BASE-T Ethernet	2	N / A	N / A	X	OSC, OSD L2/L3**, OSE, OSN *	Copper, RJ-45
3368	OSA-E2 10 GbE LR	1	N / A	01/05	X	OSD L2/L3**	Fiber, SC Duplex



LX = Long wavelength transceiver, SX = Short wavelength transceiver, LR - Long Reach transceiver
 X = Available for ordering C = Carry forward on an upgrade from z900 or z990
 * = OSN is exclusive to z9-109. Hardware availability is 09/16/05
 ** = L2/L3 = Layer 2/Layer 3 which is applicable to z9-109, z990, z890

Copper (RJ45) cabling

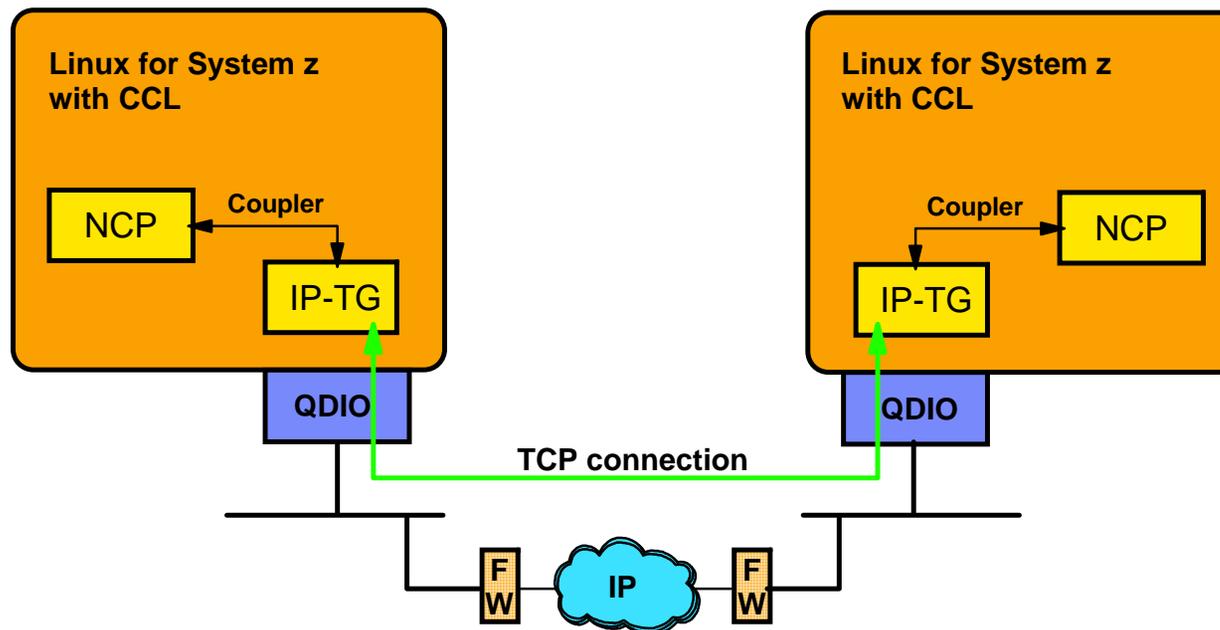
SNA LLC2 traffic over OSA LCS interfaces

- **On mainframe hardware levels that do not support QDIO layer 2 mode, an OSA Token-ring or Ethernet copper port operating in LAN Channel Station (LCS) mode can be used for SNA LLC2 access by a CCL NCP TIC2 or TIC3 interface.**
 - ▶ G5, G6, z800, z900 do not support QDIO layer 2 mode

- **OSA/SF is needed for locally administered MAC addresses on OSA LCS ports and for maintenance of the OSA Address Table (OAT) when sharing OSA LCS ports between multiple Linux images.**
 - ▶ Locally administered MAC addresses can alternatively be set via the Hardware Management Console (HMC)

- **OSA LCS ports can be shared by multiple NCP lines (in the same CCL or across multiple CCLs) - given certain restrictions, of which the most important are:**
 - ▶ For G5/G6 OSA-Express (LCS mode) and for OSA-2 in general, a port cannot be shared, but must be dedicated to the Linux image in question in TCP/IP passthru-mode.
 - ▶ For zSeries OSA-Express (LCS mode), a port can be shared by configuring unique SAP numbers in OSA/SF and the CCL NCP line definitions
 - For HOME IP address - specify an address of 0.0.0.x in OSA/SF, where 'x' is the SAP number you want assigned
 - The OSA-Express microcode must be a level 3.50 for z900 and z800 - and 5.50 for z990 and z890
 - VTAM and CCL cannot share an OSA port for communication between them - VTAM's LSA port cannot be the same as CCL's LCS port
 - Two CCL NCP line definitions cannot share an OSA LCS port for BNN traffic (they both need SAP 04)
 - One BNN NCP line definition and one or more INN/SNI NCP line definitions can share an LCS port using different local SAPs for the INN/SNI traffic

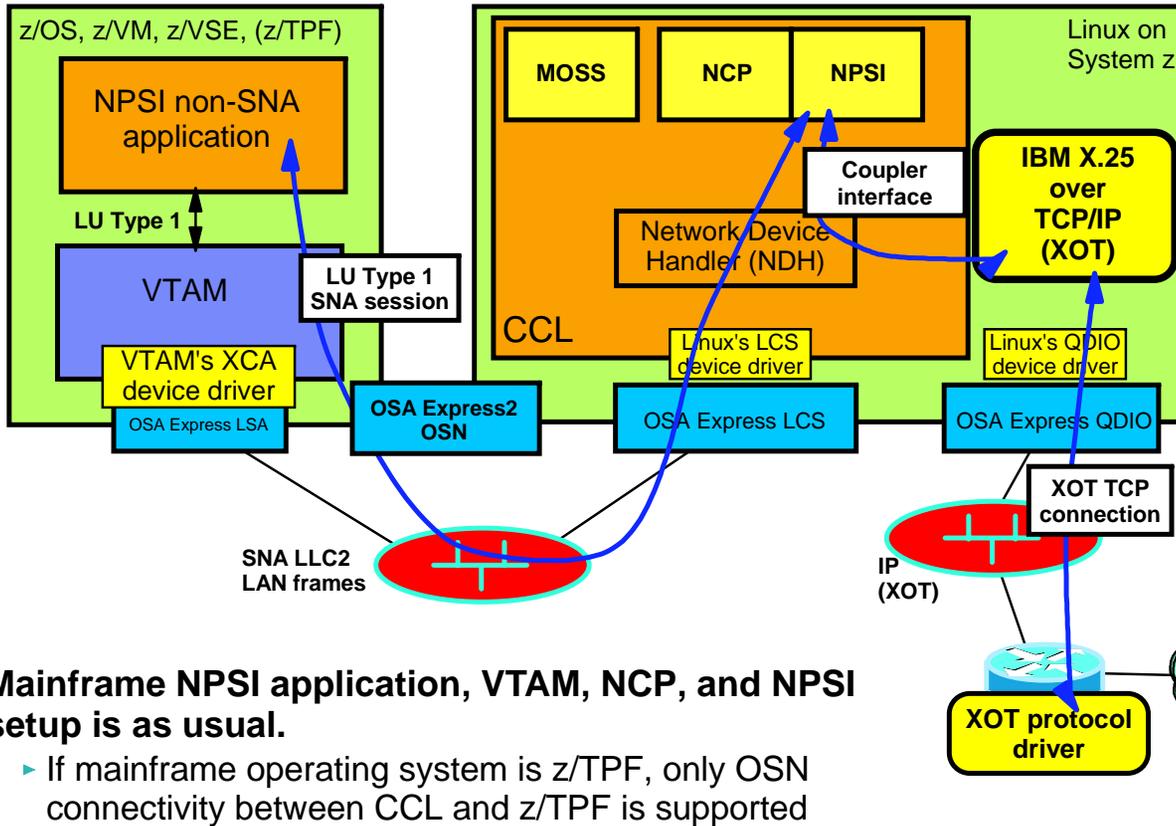
IP-TG - IP transmission group for INN/SNI traffic between two CCL NCPs



IP-TG in combination with the CDLC connectivity to VTAM provides up to 6 times better throughput than two IBM 3745 INN/SNI NCPs connected via token-ring.

- IP transmission group exchanges INN/SNI traffic between two CCL NCPs over a TCP connection.
- The NCP sees the IP-TG endpoint as a TIC3 token-ring adapter
 - ▶ TIC3 adapters normally reside in the IBM 3746 frame and are attached to Token-Ring Processors (TRP)
 - ▶ A TRP in a real IBM 3746 does all the SNA LLC2 processing on behalf of the NCP
 - No LLC2 overhead in the NCP
 - ▶ The NCP interfaces to the TRP using the coupler programming interface
 - Used by the NCP for all line and channel resources that are located in an IBM 3746
- Because there is no real LLC2 processing when using IP-TG, IP-TG performs very well for INN/SNI traffic between two CCL NCPs.
- The IP-TG TCP connection can optionally be secured (encrypted) using the STUNNEL technology of Linux.
- Configuration options allow for control of port numbers and IP addresses at for easier firewall configuration between business partners.

Transporting non-SNA X.25 access to CCL/NPSI over an IP Network - X.25 over TCP

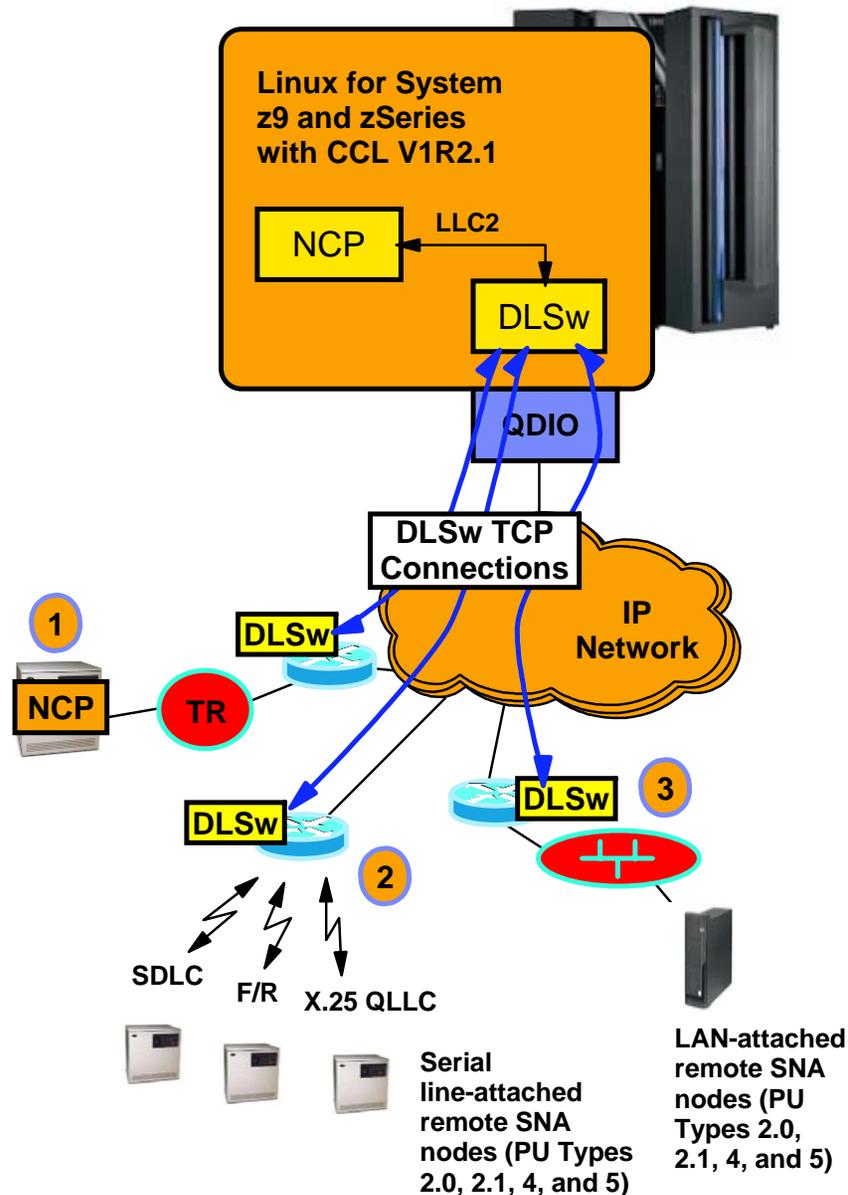


- XOT is an open standard and defined in RFC 1613 "Cisco Systems X.25 over TCP (XOT)".
- XOT is used to encapsulate X.25 packets over a TCP/IP network.
 - ▶ Supported by various router vendors - including Cisco

XOT with CCL enables continued use of NPSI-based applications and X.25 connectivity to the System z platform via an XOT router.

- Mainframe NPSI application, VTAM, NCP, and NPSI setup is as usual.
 - ▶ If mainframe operating system is z/TPF, only OSN connectivity between CCL and z/TPF is supported
- NPSI processing remains offloaded from the mainframe OS environment.
- Physical connectivity to X.25 network is via an aggregation layer router.
 - ▶ Connectivity between aggregation layer router and NPSI is via an X.25 Over TCP/IP (XOT) TCP connection (IP network flows).
- Interface between NPSI and local XOT protocol component is the same as NPSI uses today when communicating over X.25 adapters in an IBM 3746 unit - the Coupler interface.
- X.25 over TCP/IP for CCL is a separate IBM software product that is needed in conjunction with CCL for X.25 connectivity to NPSI

CCL imbedded DLSw support



➤ Network infrastructure simplification:

- ▶ Integration of data center DLSw functions with NCP functions in Linux on System z
- ▶ Avoids or reduces the need for separate data center DLSw router equipment

➤ CCL DLSw is based on the open standards version of DLSw - RFC1795 & RFC2166:

- ▶ Interoperability with all vendors who have implemented DLSw according to those standards
- ▶ Will interoperate with DLSw+ nodes
 - DLSw+ nodes will adapt to standard DLSw protocols when connecting to an open standards-based DLSw implementation

➤ Typical CCL DLSw scenarios:

1. INN/SNI to NCPs in remote IBM 3745/46
2. Peripheral nodes attached via serial lines to DLSw router
3. Peripheral nodes attached via LAN to DLSw router

➤ Virtualizes the SNA MAC address

- ▶ A DLSw MAC address is a virtual MAC address that doesn't need to match any real MAC address
- ▶ Supports many virtual DLSw MAC addresses over a single physical MAC address (OSA port)
 - Especially of value on mainframe hardware platforms where QDIO layer-2 mode isn't available

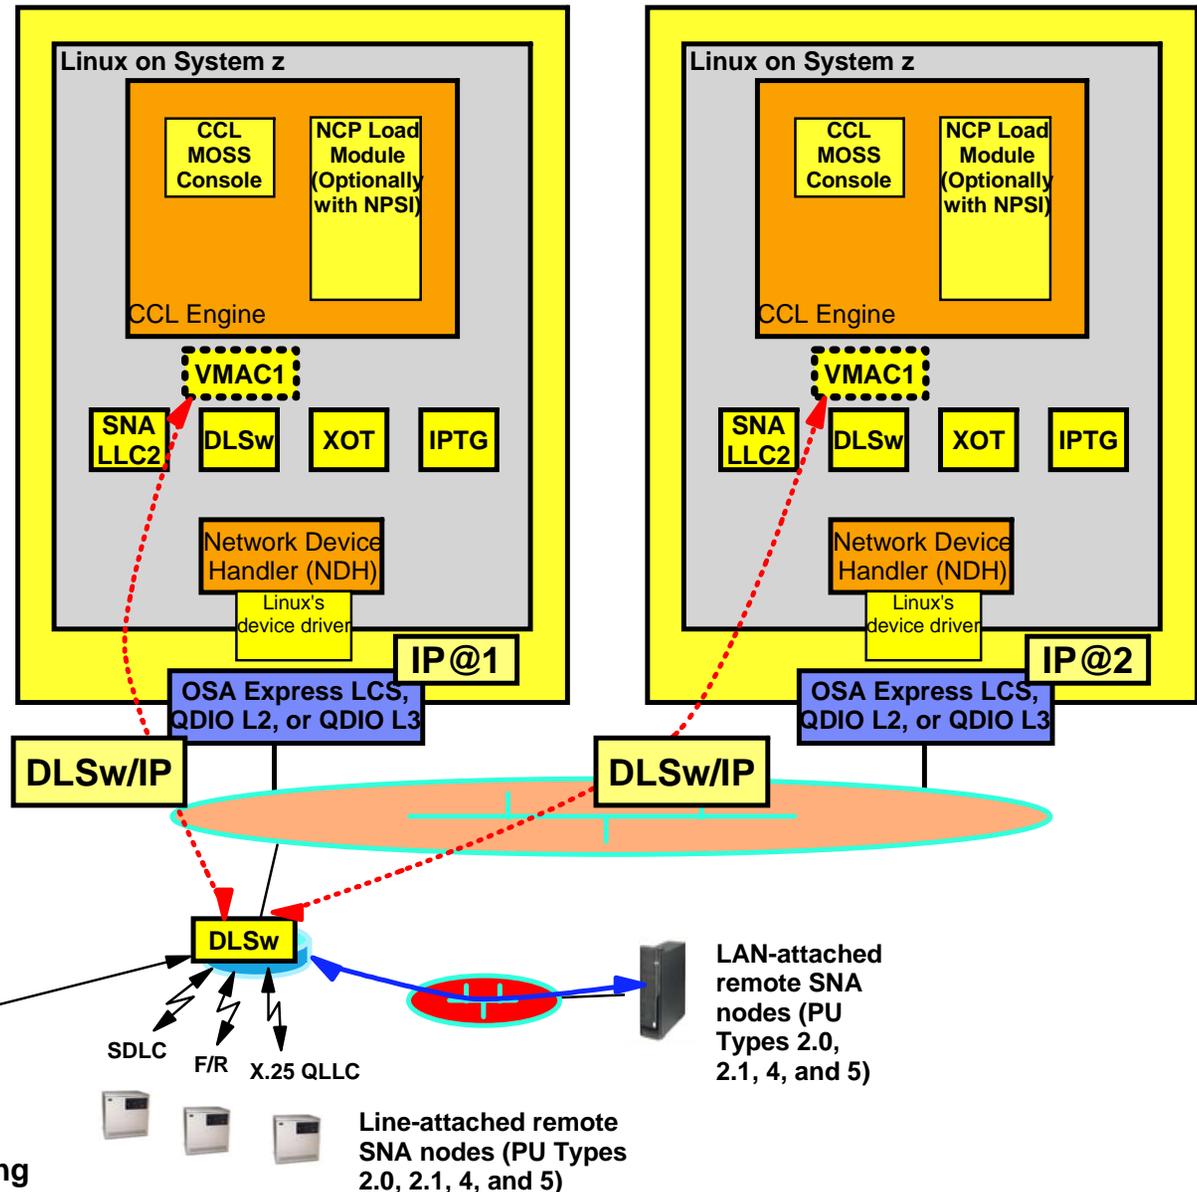
Duplicate MAC addressing for load-balancing and redundancy

➤ Duplicate TIC addressing is a token-ring only technology

- ▶ Works also for CCL if OSA token-ring ports are used
- ▶ It is not supported by an Ethernet LAN infrastructure
- ▶ Duplicate TIC address topologies cannot be directly migrated to Ethernet

➤ The same objectives can be addressed by using a DLSw infrastructure

- ▶ "Remote" DLSw node defines two CCL NCP DLSw peers
- ▶ Both peers support the same SNA MAC address
 - Both CCL/NCPs have a token-ring LINE with that same MAC address defined
- ▶ The remote DLSw node can be configured to balance LLC2 connections over the two peers when both are available
- ▶ Simple round-robin and circuit-weighted round-robin are typically supported

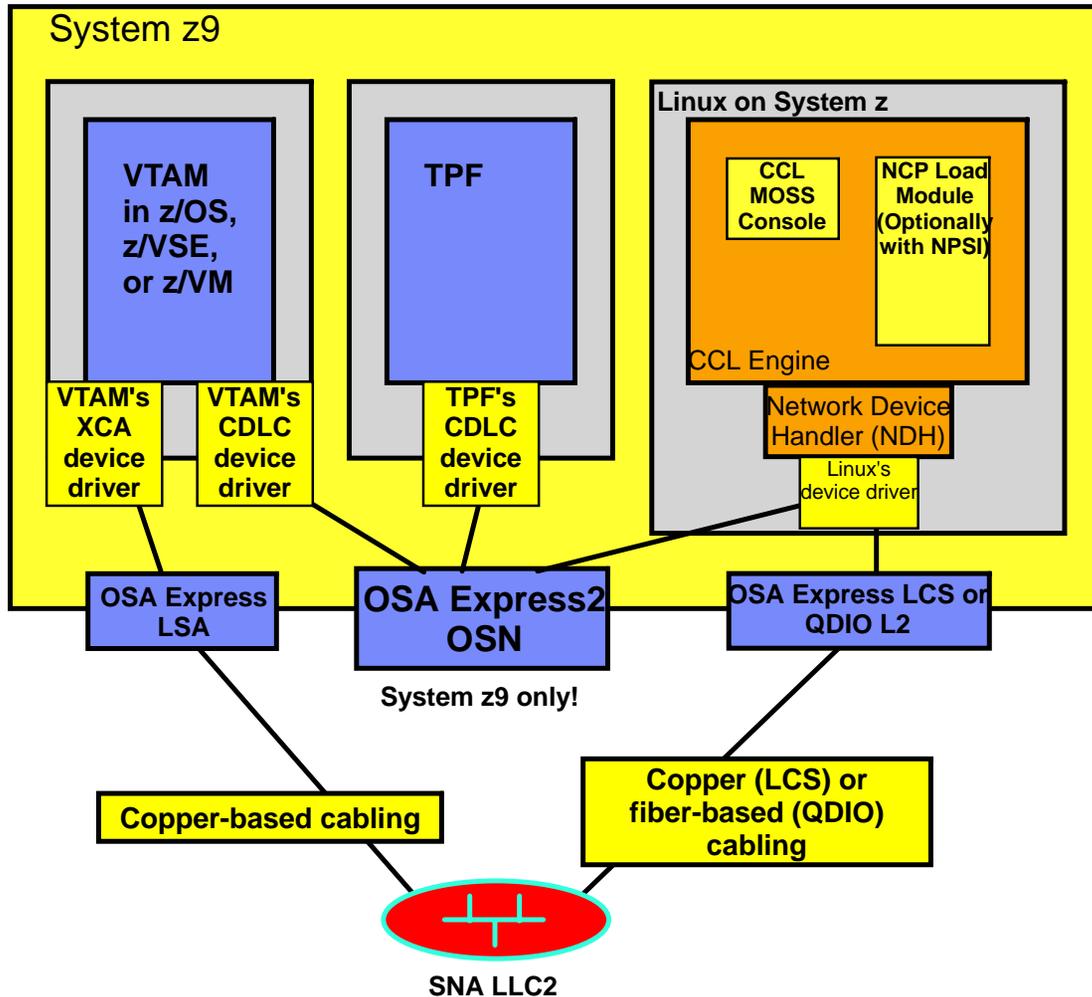


Two remote peers:

- ▶ IP@1
- ▶ IP@2

Load-balancing either round-robin or circuit weighting

CCL and VTAM/TPF connectivity summary

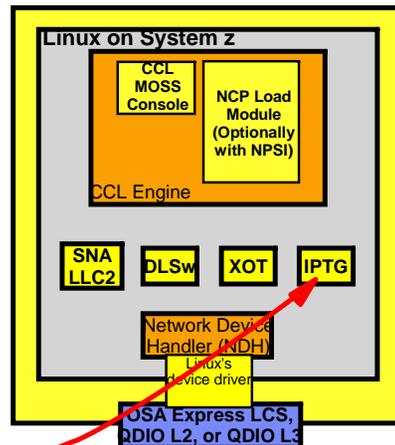
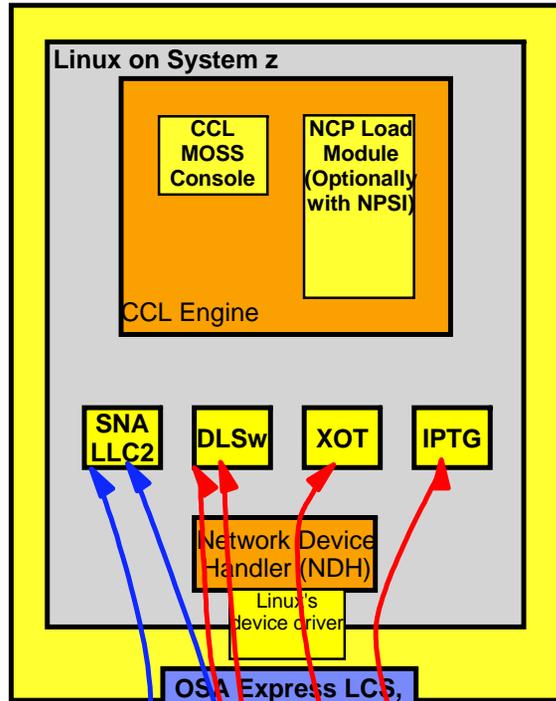


Note: The OSN technology is available on System z9 hardware only.

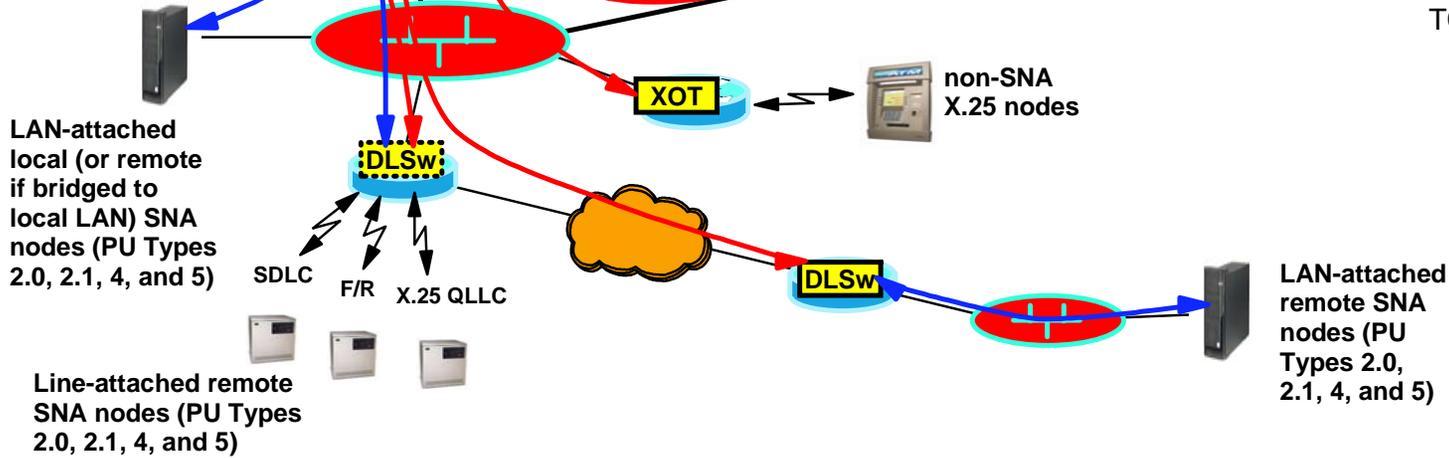
- **VTAM connects to a CCL NCP using one of two technologies:**
 - ▶ Over a LAN to which VTAM connects using an OSA port in LSA mode and Linux over an OSA port in LCS or QETH mode (QDIO Layer 2).
 - ▶ If VTAM and CCL reside on the same System z9, they can connect via a shared OSA-E2 port operating in OSA for NCP (OSN) mode.
 - Both VTAM, TPF, and the NCP see this connectivity as an ESCON channel over which the usual CDLC channel protocol is used.
- **TPF supports the OSN connectivity option only.**
- **No SNA subarea topology changes - VTAM is still a PU Type 5 and the NCP is a PU Type 4**
 - ▶ In most cases no changes to SNA subarea pathing definitions
- **When OSN connectivity is used, there are no changes to VTAM definitions or VTAM operations procedures.**
- **When LAN connectivity is used, there may be minor changes to VTAM definitions and VTAM operations procedures.**
- **In most cases no changes to NetView definitions and operations.**

CCL and down-stream connectivity summary

OSA mode used by CCL	Networking protocols	Cabling
LCS	SNA LLC2 and IP	Copper, RJ45
QDIO Layer 2	SNA LLC2 and IP	Fiber or copper
QDIO Layer 3	IP	Fiber or copper



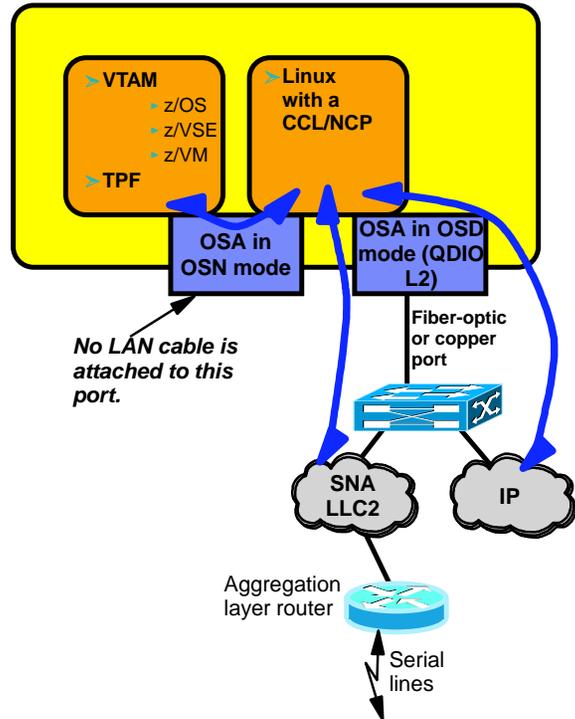
- > **SNA LLC2**
 - ▶ SNA protocol data units (PDUs) in native frames on the LAN
- > **DLSw**
 - ▶ SNA Subarea and/or APPN ISR PDUs encapsulated in a TCP connection
 - ▶ Local DLSw refers to serial line termination and PDU switching to a LAN
- > **XOT**
 - ▶ X.25 packets encapsulated in a TCP connection
- > **IPTG**
 - ▶ SNA INN/SNI PDUs encapsulated in a TCP connection



The NCP sees all down-stream resources as though they were Token-ring attached - NCP TIC2 or TIC3 line address.

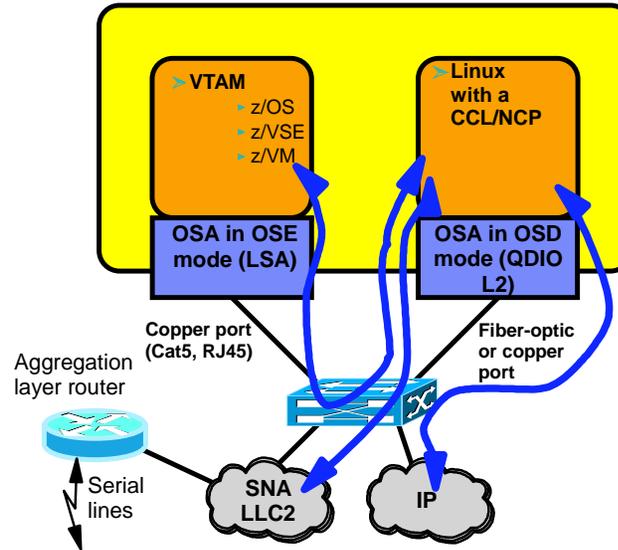
Mainframe hardware level and CCL connectivity summary

System z9



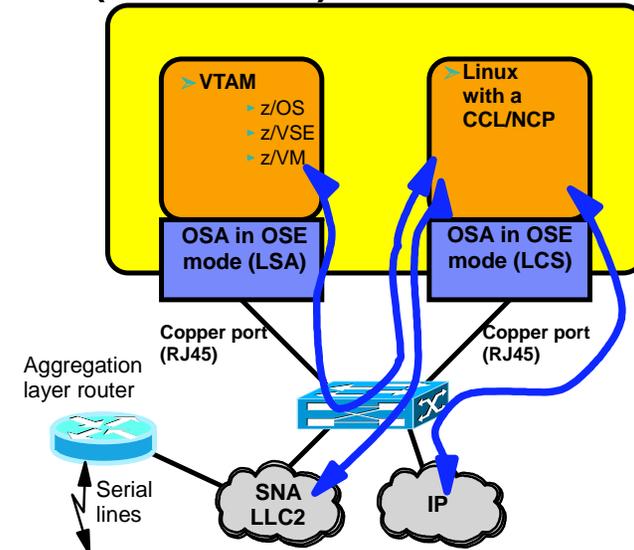
- **Owning VTAM**
 - OSN (CDLC)
 - Via shared LAN
- **SNA LLC2**
 - QDIO layer 2
 - LCS
- **IP**
 - IP-TG
 - DLSw
 - XOT

zSeries z890/z990



- **Owning VTAM**
 - Via shared LAN
- **SNA LLC2**
 - QDIO layer 2
 - LCS
- **IP**
 - IP-TG
 - DLSw
 - XOT

zSeries z800/z900 (and G5/G6)

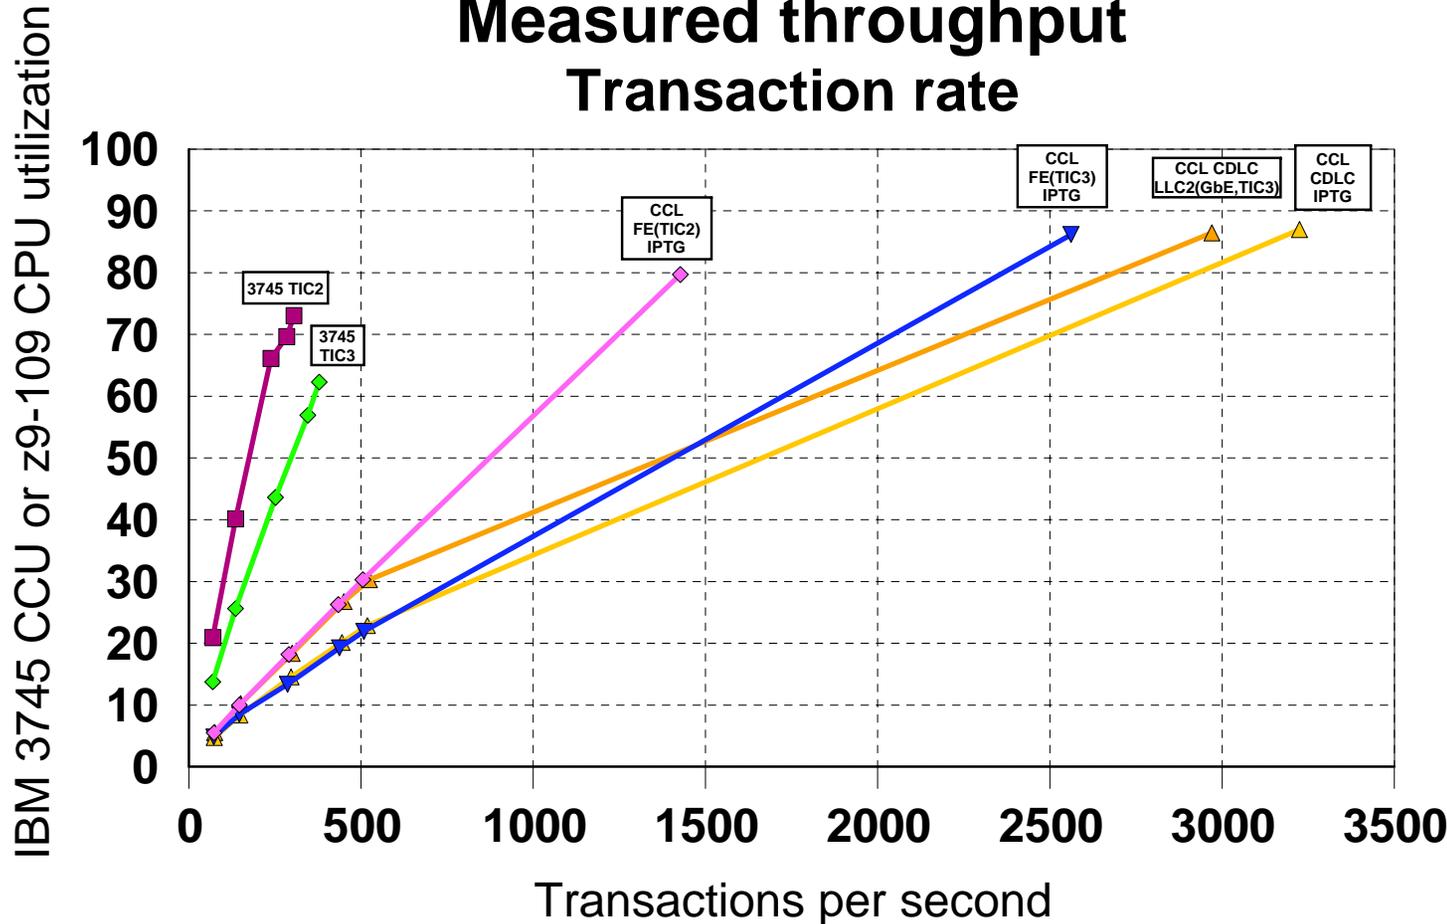


- **Owning VTAM**
 - Via shared LAN
- **SNA LLC2**
 - LCS
- **IP**
 - IP-TG
 - DLSw
 - XOT

Performance and Capacity

SNI performance - throughput

SNI Transactional Workload Measured throughput Transaction rate



- > **FE(TIC2),IPTG**
 - ▶ Fast Ethernet between VTAMs and CCL/NCPs using TIC2 ports and IPTG between SNI CCL/NCPs
- > **FE(TIC3),IPTG**
 - ▶ Fast Ethernet between VTAMs and CCL/NCPs using TIC3 ports and IPTG between SNI CCL/NCPs
- > **CDLC,LLC2(GbE,TIC3)**
 - ▶ OSN between VTAMs and CCL/NCPs and SNA LLC2 over QDIO layer2 via a Gigabit Ethernet between SNI CCL/NCPs
- > **CDLC,IPTG**
 - ▶ OSN between VTAMs and CCL/NCPs and IPTG between SNI CCL/NCPs

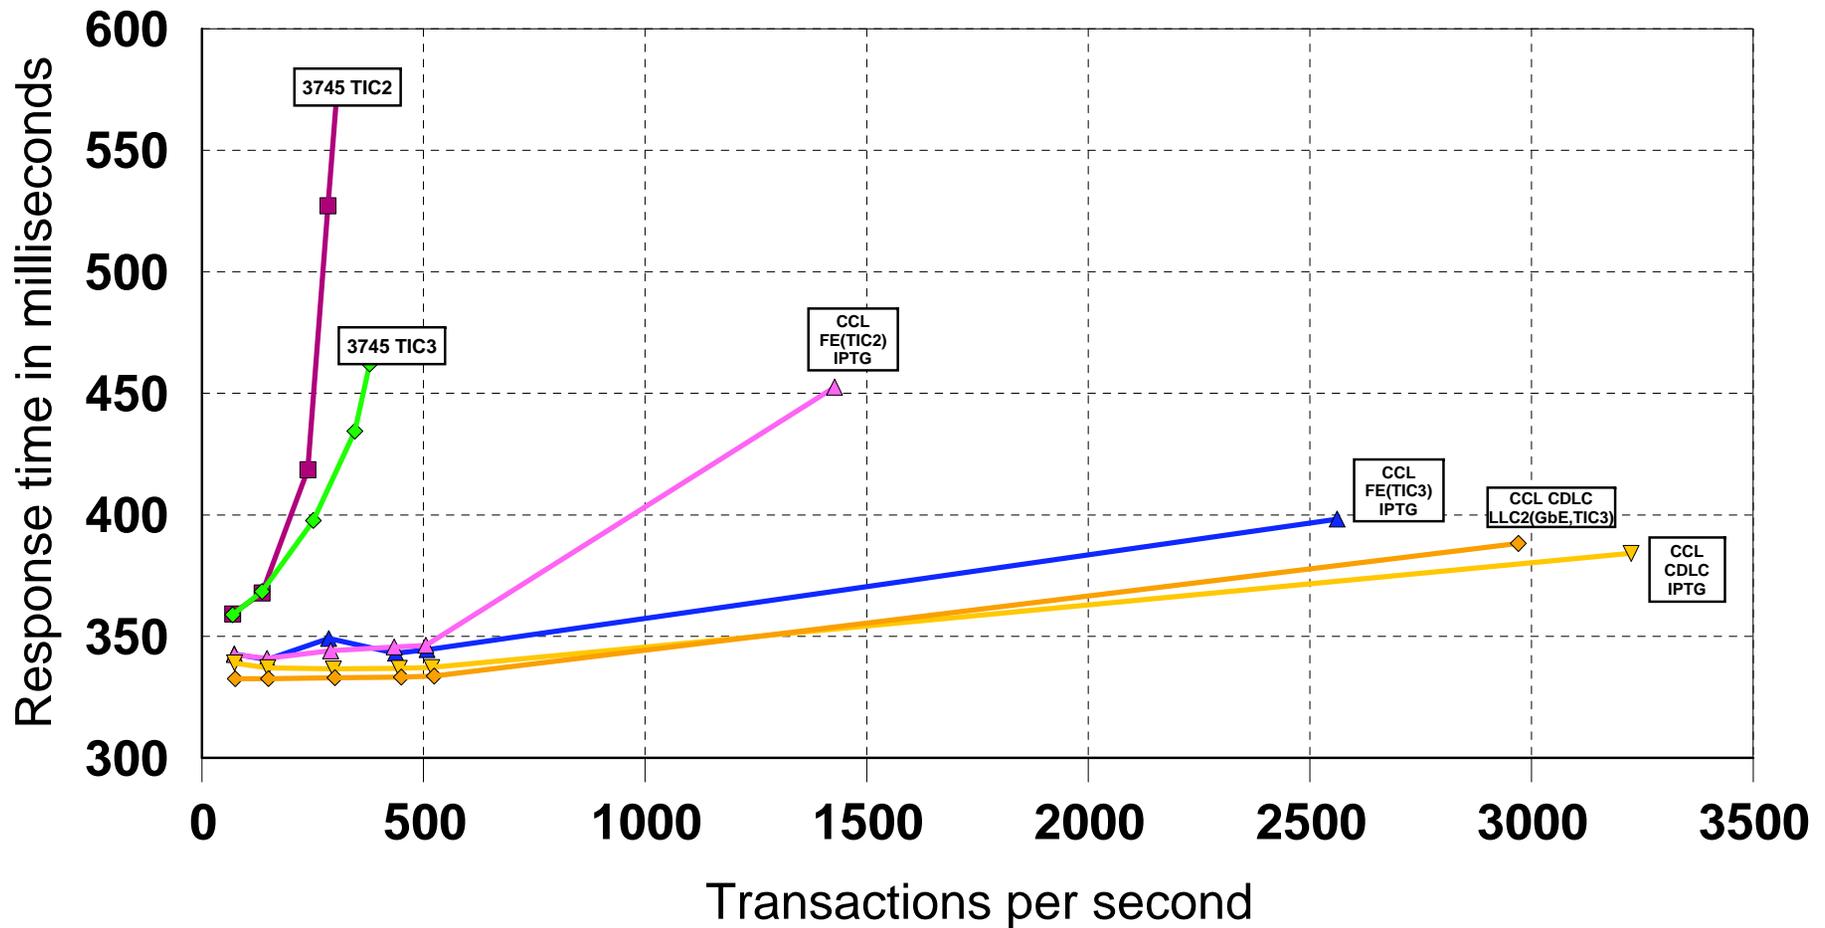
An IBM 3745-31A sample configuration with TIC3 ports maxes out around 380 transactions per second

Note: A wide area network in-between the two NCPs will have an impact on transaction rate.

Note: All CCL data was captured with CCL V1.2.1

SNI performance - response time

SNI Transactional Workload Measured response time



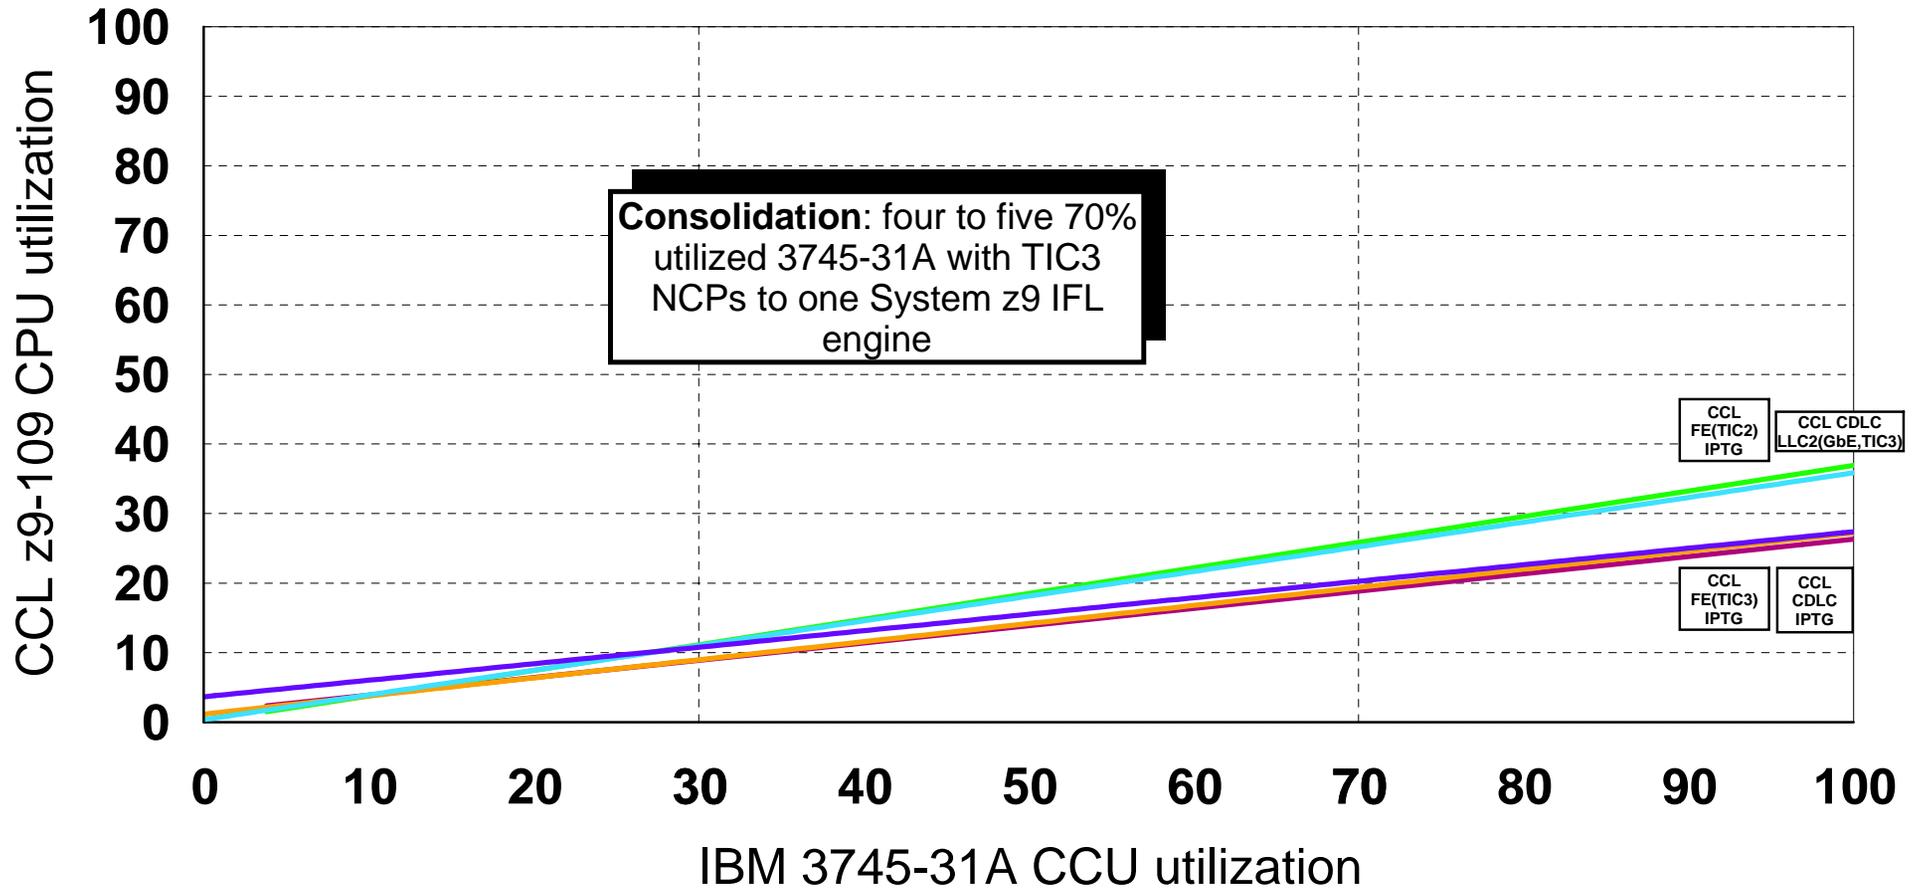
Note: The response time shown above includes an artificial thinktime between transactions of 330 milliseconds.

Note: The high-transaction volume measurement points were performed with a think time of zero - they have in this chart been increased with 330 milliseconds for comparison reasons.

SNI workload CPU capacity planning data for CCL V1R2.1 running on System z9

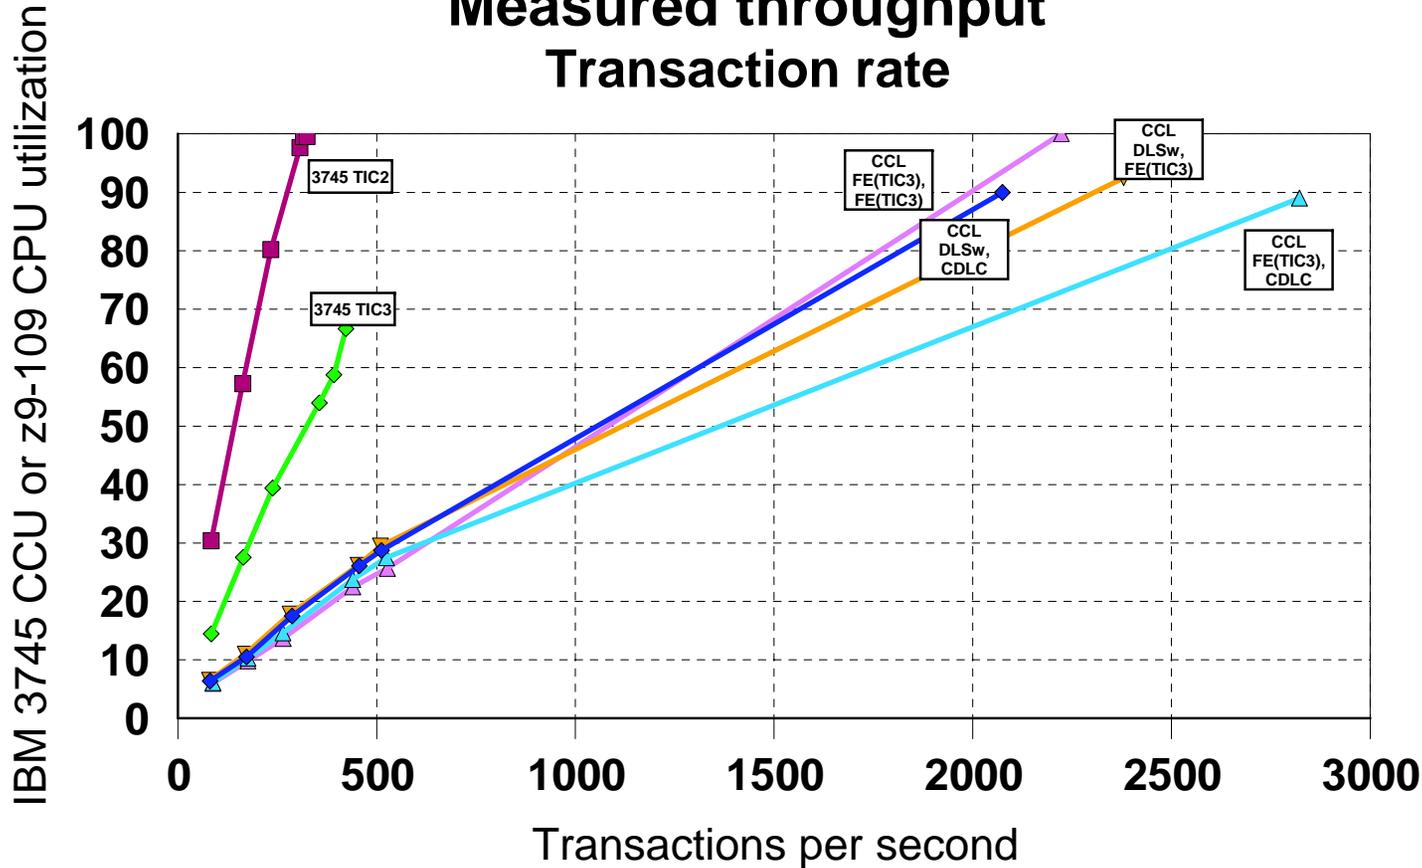
SNI Transactional Workload

IBM 3745-31A CCU utilization vs. z9-109 CPU utilization IBM 3745-31A with TIC3 adapters



Boundary function performance - throughput

Boundary Function Transactional Workload Measured throughput Transaction rate



- **FE(TIC3),FE(TIC3)**
 - ▶ Peripheral nodes SNA LLC2 via Fast Ethernet, CCL/NCP using TIC3 both downstream and upstream to VTAM
- **FE(TIC3),CDLC**
 - ▶ Peripheral nodes SNA LLC2 via Fast Ethernet, CCL/NCP using TIC3 downstream and CDLC upstream to VTAM
- **DLSw,FE(TIC3)**
 - ▶ Peripheral nodes via DLSw to CCL/NCP using TIC3 upstream to VTAM
- **DLSw,CDLC**
 - ▶ Peripheral nodes via DLSw to CCL/NCP using CDLC upstream to VTAM

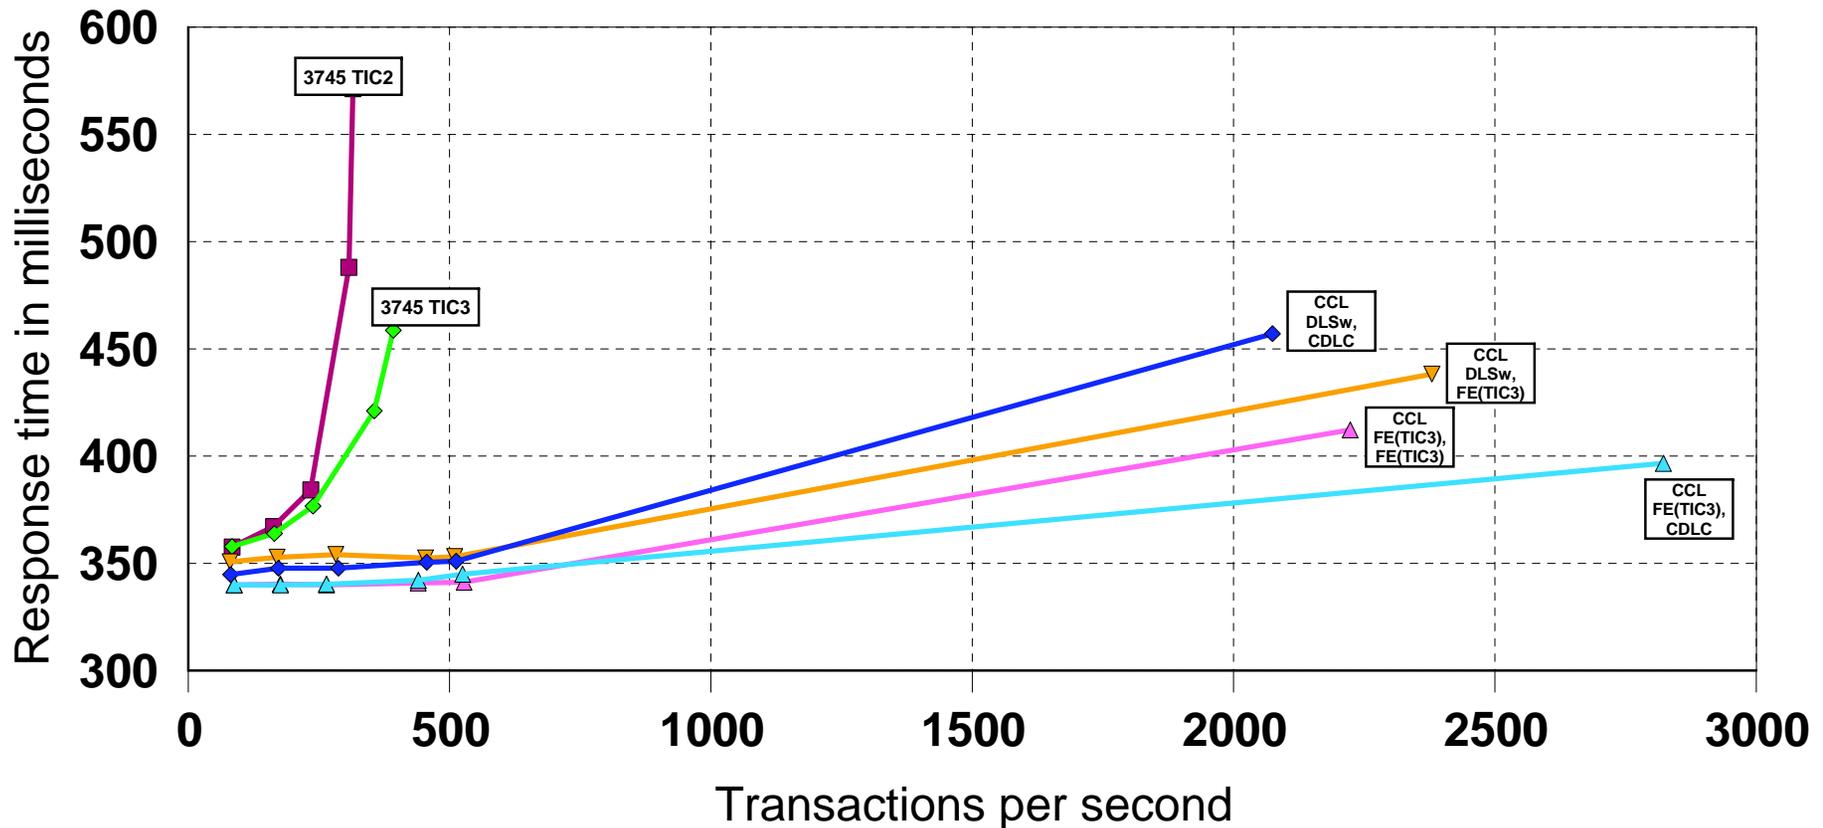
An IBM 3745-31A sample configuration with TIC3 ports maxes out around 268 transactions per second

Note: Clients emulated by four Linux servers. For DLSw workloads, clients are connected via two Cisco 7507 DLSw routers over GbE to CCL's imbedded DLSw component.

Note: All CCL data was captured with CCL V1.2.1

Boundary function performance - response time

Boundary Function Transactional Workload Measured response times

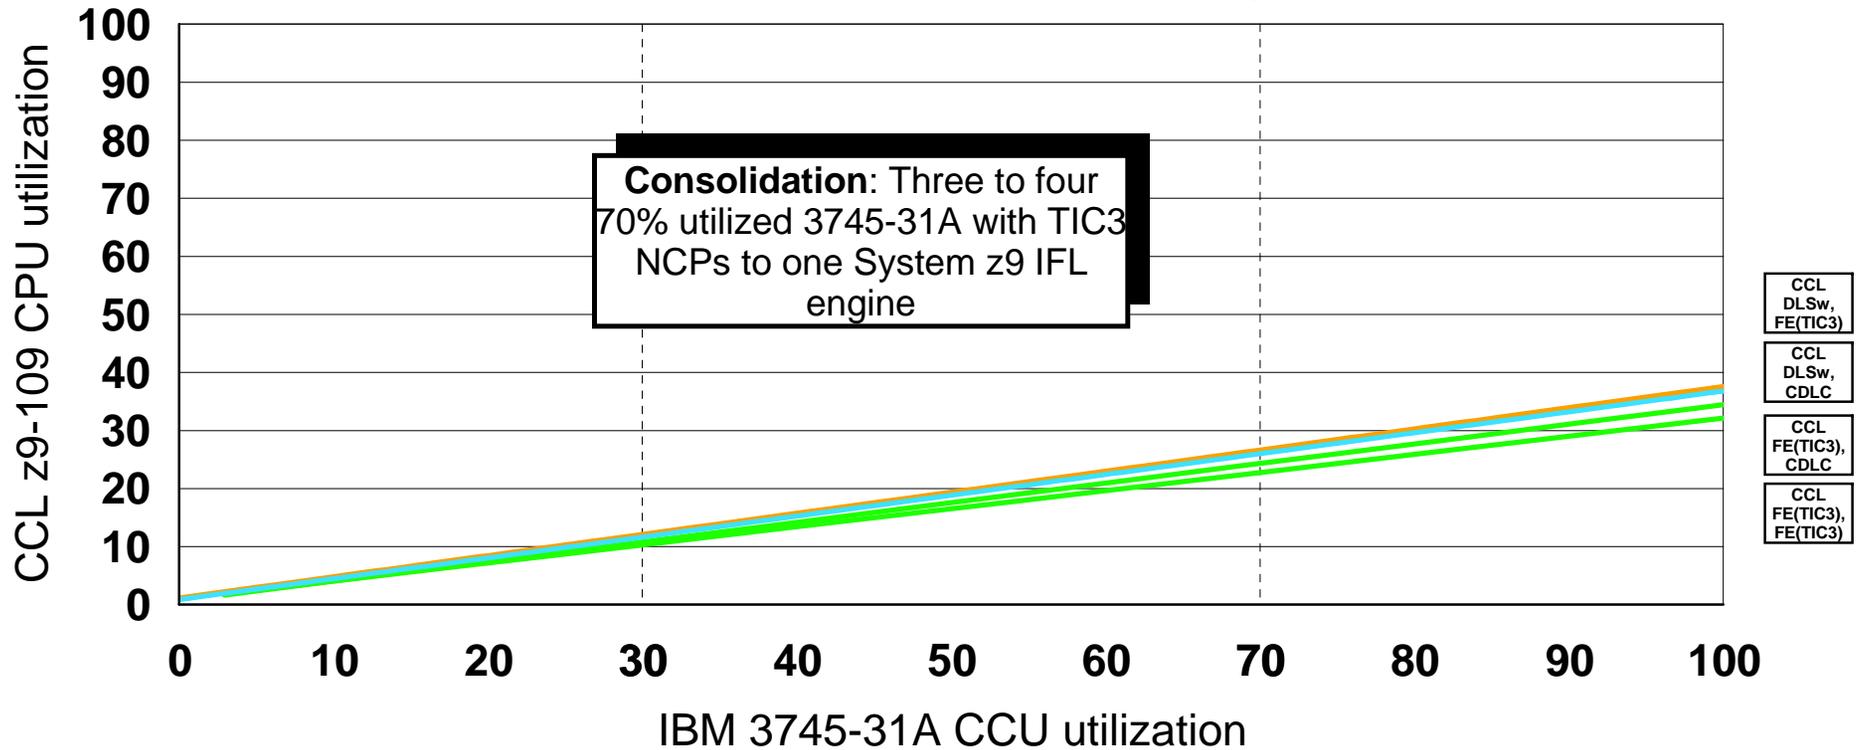


Note: The response time shown above includes an artificial thinktime between transactions of 330 milliseconds.

Note: The high-transaction volume measurement points were performed with a think time of zero - they have in this chart been increased with 330 milliseconds for comparison reasons.

Boundary function workload CPU capacity planning data for CCL V1R2.1 running on System z9

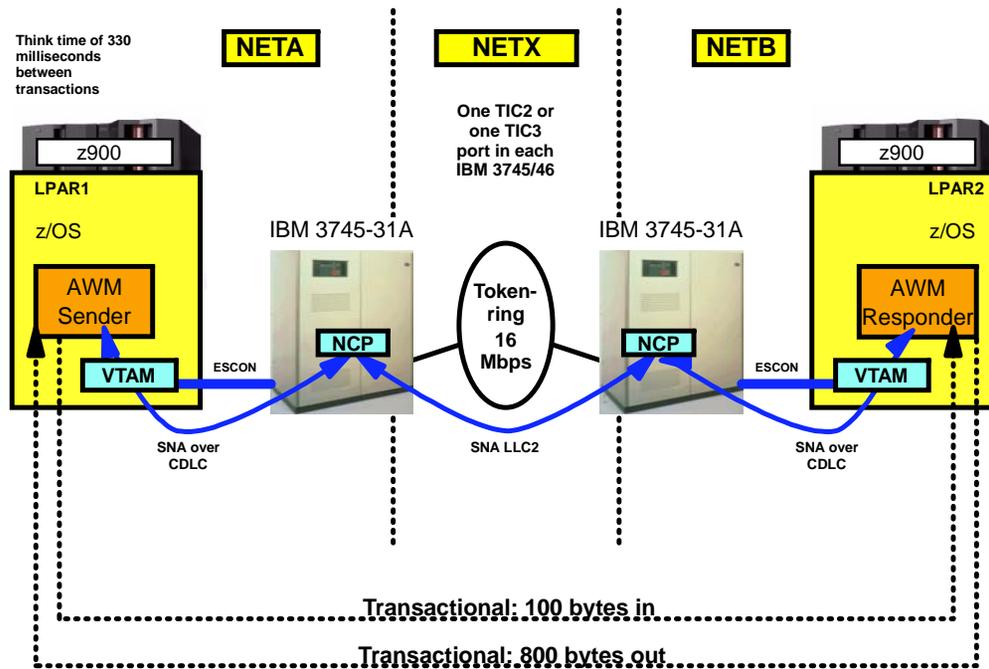
Boundary Function Transactional Workload IBM 3745-31A CCU utilization vs. z9-109 CPU utilization IBM 3745-31A with TIC3 adapters



CCL CPU usage - SNI transactional workload - setup notes

NOTES

IBM 3745/46 test environment overview



➤ **IBM 3745-61A**

- ▶ Tests were done with an IBM 3745-61A - divided into two IBM 3745-31A units
- ▶ CCU capacity equals an IBM 3745-31A
- ▶ One TIC2 or one TIC3 adapter

➤ **CCL as z/VM guest**

- ▶ 2 dedicated z990 or z9 CPs to z/VM
- ▶ 1 virtual CP to each CCL Linux guest
- ▶ "Client" side CCL working set 105 MB (as reported by z/VM)
- ▶ "Server" side CCL working set 340 MB (as reported by z/VM)

➤ **z990 or z9-109 hardware for CCL**

➤ **CCL V1.2.1: SUSE SLES9 SP2 - Linux 2.6 kernel**

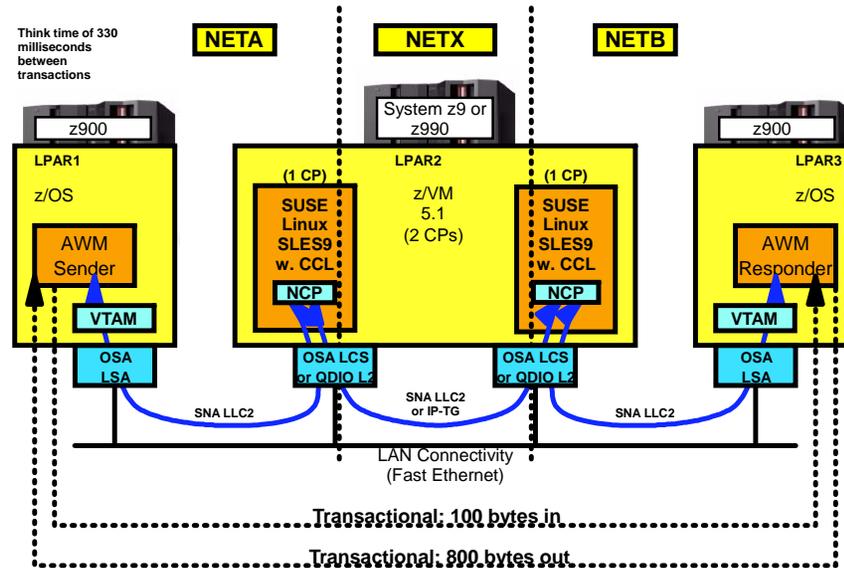
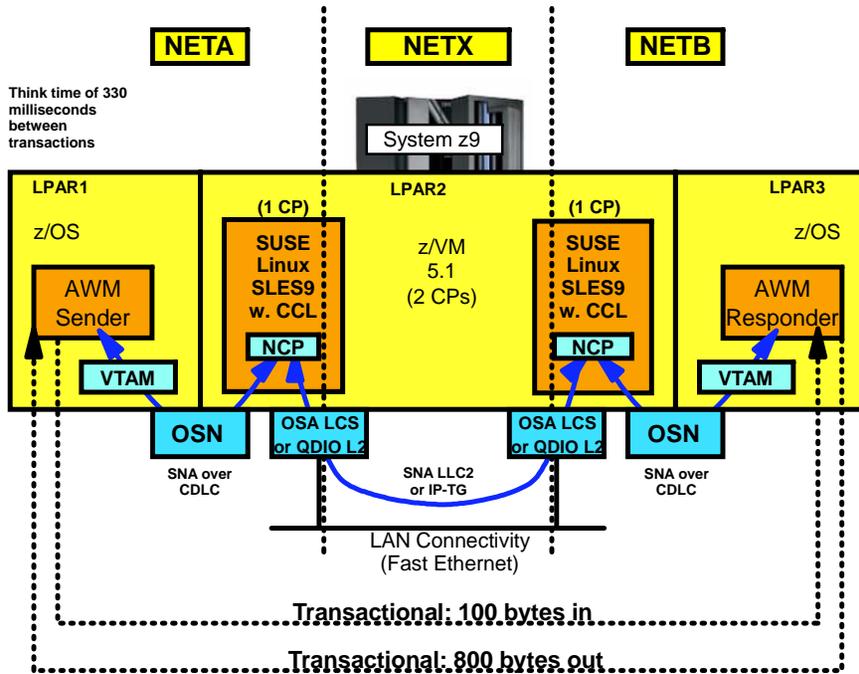
➤ **Transactional workload characteristics**

- ▶ Up to 175 LU 6.2 sessions used to generate traffic
- ▶ 100 bytes in per transaction
- ▶ 800 bytes out per transaction
- ▶ Thinktime 330 milliseconds between transactions per SNA session
 - High transaction volume measurements achieved with a thinktime of zero
- ▶ VTAM IOBUF size=932
- ▶ IBM 3745 MAXBFRU=20
- ▶ All workloads were driven by AWM

CCL CPU usage - SNI transactional workload - setup notes

NOTES

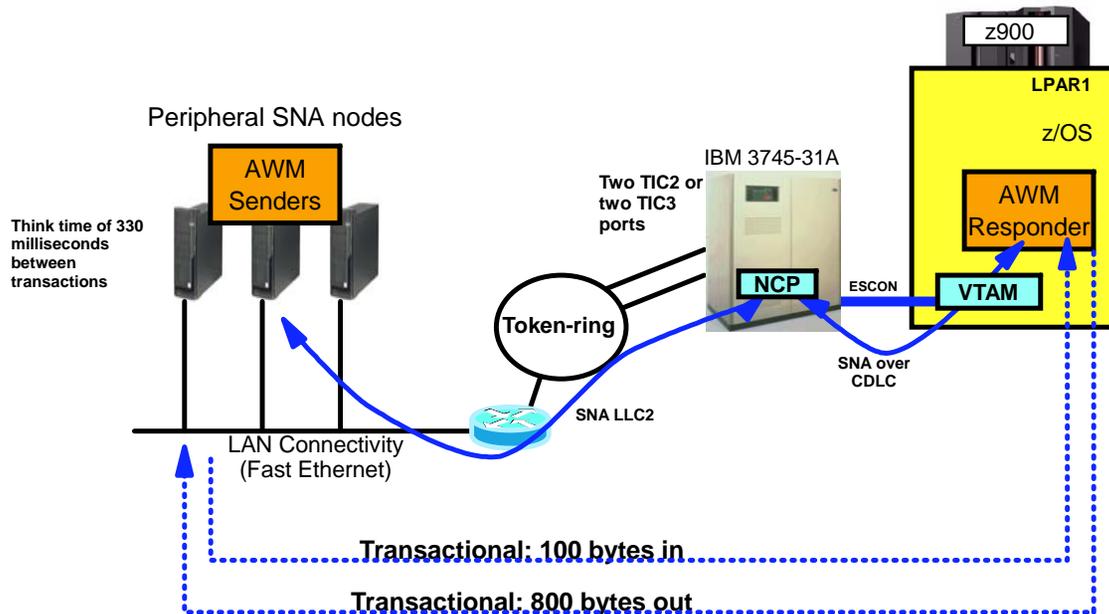
IBM CCL using shared LAN between VTAM and the CCL NCP:



IBM CCL using OSA for NCP (OSN) CDLC connectivity between VTAM and the CCL NCP:

CCL CPU usage - boundary function transactional workload - setup notes

IBM 3745/46 test environment overview



> **IBM 3745-31A**

- ▶ CCU capacity equals an IBM 3745-31A
- ▶ Two TIC2 or two TIC3 adapters

> **CCL as z/VM guest**

- ▶ 2 dedicated z990 or z9 CPs to z/VM
- ▶ 1 virtual CP to the CCL Linux guest

> **z990 or z9-109 hardware for CCL**

> **CCL V1.2.1: SUSE SLES9 SP2 - Linux 2.6 kernel**

> **Transactional workload characteristics**

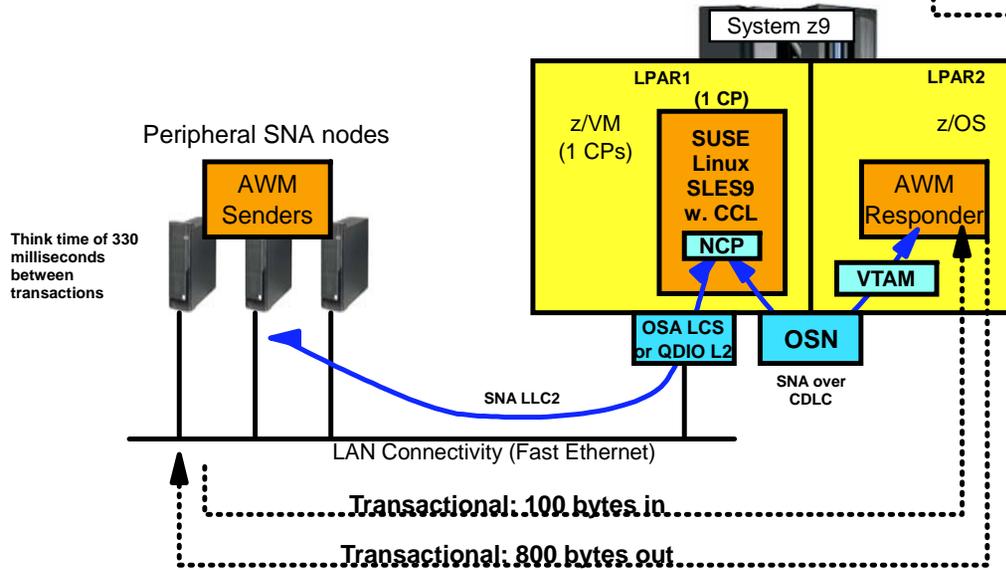
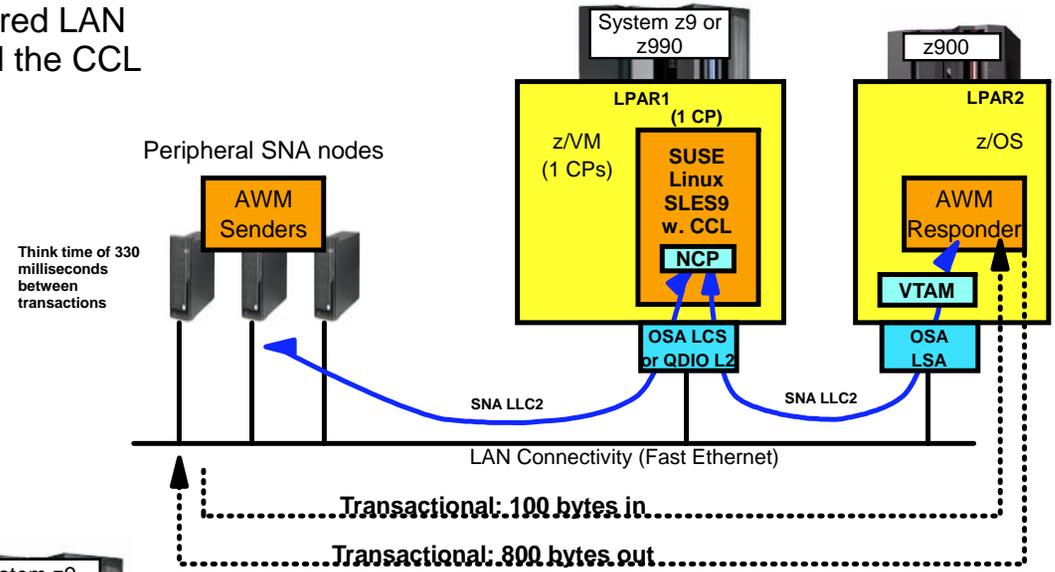
- ▶ BF devices were emulated as one LU per PU SNA devices
- ▶ Each transaction consisted of a 100-byte request in and a 800-byte response
- ▶ A thinktime of 330 milli seconds was used between each transaction
 - High transaction volume measurements achieved with a thinktime of zero
- ▶ VTAM IOBUF size 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.
- ▶ All workloads were driven by AWM

NOTES

CCL CPU usage - boundary function transactional workload - setup notes

NOTES

IBM CCL using shared LAN between VTAM and the CCL NCP:

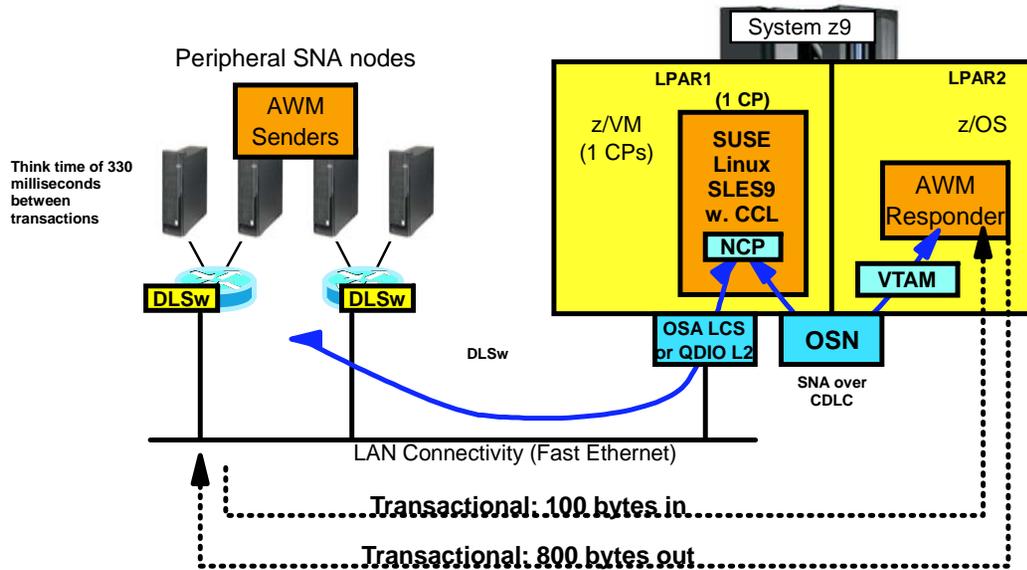
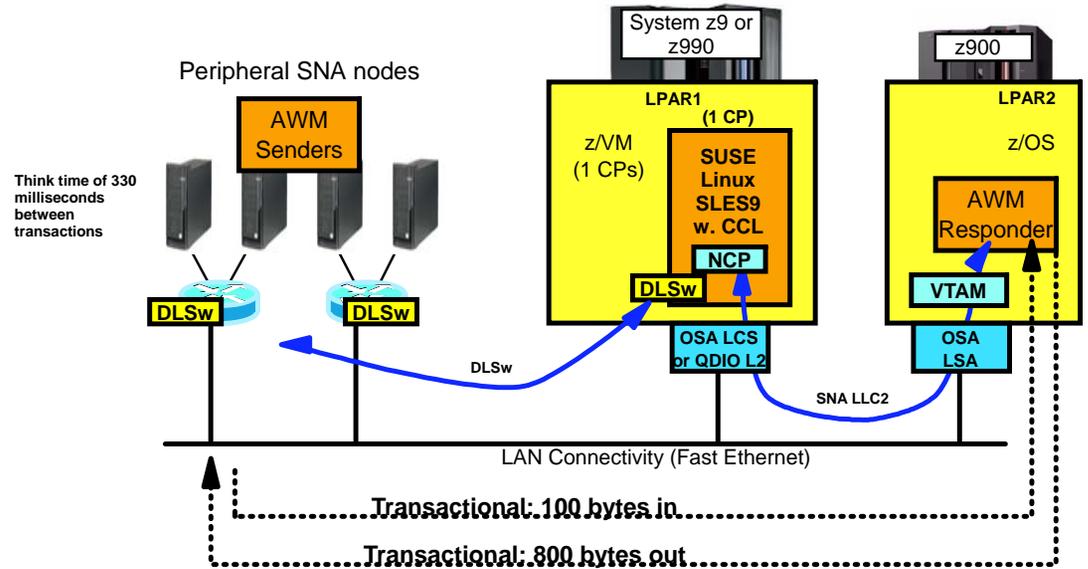


IBM CCL using OSA for NCP (OSN) CDLC connectivity between VTAM and the CCL NCP:

CCL CPU usage - boundary function transactional workload - setup notes

NOTES

DLSw between peripheral nodes and the CCL NCP and shared LAN between VTAM and the CCL NCP:



DLSw between peripheral nodes and the CCL NCP and OSA for NCP (OSN) CDLC connectivity between VTAM and the CCL NCP:

Summary

CCL offers an opportunity to simplify and improve the SNA hardware infrastructure

- **Removes the IBM 3745/46 hardware component for most of the current usage scenarios:**
 - ▶ Reduced need for raised floor space
 - ▶ Reduced requirements for power and cooling capacity
 - ▶ Reduced need for skills and resources to manage the physical IBM 3745/46 hardware

- **Reduces the requirement for continued use of ESCON technology in the data center:**
 - ▶ Removes the need for System z ESCON channel interfaces for NCP connectivity
 - ▶ Removes the need for ESCON director ports for NCP connectivity
 - ▶ If IBM 3745/46 is the last hardware component that uses ESCON directors, removes the ESCON directors themselves – including the skills and resources associated with managing the ESCON directors

- **A CCL NCP is not limited to token-ring LANs, but can use any LAN technology that is supported by an OSA port in OSE (LCS) or OSD (QDIO) mode:**
 - ▶ Includes token-ring on System z platforms where OSA continues to support token-ring
 - ▶ 10/100/1000 Mbps 10BASE-T Ethernet – cat5 cabling, RJ45
 - ▶ 1 Gbps and 10 Gbps Ethernet – fiber optic cabling
 - ▶ Removes the need for continued use of token-ring LAN equipment:
 - Access units
 - Token-ring ports in switch equipment
 - Token-ring LAN management skills and software

- **CCL through LAN interface virtualization, provides a much more efficient use of overall LAN capacity:**
 - ▶ QDIO layer 2: up to 2048 virtual LAN interfaces on one OSA port
 - ▶ DLSw virtualizes the LAN interface in a DLSw environment

CCL offers improved opportunities for implementing high-availability NCP topologies

- **CCL runs on System z and inherits all the unique availability features of the System z platform**
- **Deploying redundant stand-by NCPs can be done without acquiring additional hardware CCUs or ESCON channel hardware**
 - ▶ A stand-by dormant CCL NCP on System z uses very few resources while it is not being used
 - ▶ Much easier and cheaper to design and deploy redundancy for high availability
- **Even though Ethernet does not support duplicate Medium Access Control (MAC) addresses, a traditional SNA duplicate MAC addressing topology can be deployed using DLSw technologies:**
 - ▶ To serve as a load balancing technology where peripheral nodes contact one of more boundary function NCPs or one of more OSA ports into a single boundary function NCP
 - ▶ To server as an availability technology, where a peripheral node re-connects to the same MAC address after a failure – but now connects to a back-up NCP

CCL offers both existing and enhanced management capabilities

➤ Existing SNA management tools continue to work with an NCP running in CCL:

- ▶ SSCP takeover procedures in a traditional dual-CMC network host environment
- ▶ XRF for CICS and IMS session recovery
- ▶ Tivoli NetView for z/OS
- ▶ NTuneMon
- ▶ NPA/NPM
- ▶ OEM products

➤ The Linux platform offers management opportunities that are new to the NCP workload:

- ▶ Linux is a general-purpose open operating system with many management features built into it
- ▶ CCL runs as traditional Linux processes and can be monitored and managed as any normal Linux process
- ▶ Monitoring and automated recovery of a CCL process (a virtual CCU)
 - Restart in-place
 - Restart in other Linux instance
- ▶ Monitoring and automated actions to CCL event messages that in a real IBM 3745/46 environment went to the MOSS console
 - Messages can be consolidated to central message automation point, such as NetView on z/OS
- ▶ The structuring of CCL, NDH, and the Linux device drivers provide for more detailed insight into the flow of data between the NCP and the real network adapters
 - NDH diagnostics commands
 - Trace details that are unique to this environment
- ▶ Close integration of Linux management, CCL management, and CCL NCP management can be done using Tivoli System Automation for Operations

NTuneMon V3R2 example when used with a CCL/NCP

```

ATUSS  E74SVT6 Summary Status CCU= 0% Storage= 2% NTuneMON V3R2 11:58
-----
GENERATION INFORMATION          3745 HARDWARE INFO          SNI INFORMATION
07/08/2005 09:35:21           MICROCODE EC = CCLV1R2       SNI NETWORKS= 2
3745-31A 16MB E74SVT6         FIX = 07-12-05             HSCBS IN USE= 2400 48%
SA 74 S/N= 00000000         CDS Update= 10/07/2004    NATIVE NETID= NETE
ACF SSP V4R8.1 MVS                                     BUFFER POOL INFORMATION
ACF NCP V7R8.1F CCU A       3746 M900 INFORMATION      BUFFERS 1%
564806300 SINGLE CCU        S/N= 00666666             BPOOL 0%
USAGE TIER = 5                                                     DYNPOOL 0%
DISK LOADED NOT VTAM
-----
VR INFORMATION                 VRTP NETWORK SUBAREA CUR MAX MIN RCVDQ XMTQ VR STATE
ACTIVE/MAX= 6/103             0.2 NETE 8 17 255 15 1 0
ALARM/WARN= 0/3              0.0 NETX 78 80 255 80 0 0
USER ALARM= 100%             0.0 NETE 8 15 255 15 0 0
-----
TG INFORMATION                 TGN NETWORK SUBAREA LOW MEDIUM HIGH TOTAL
ACTIVE/MAX= 3/6
ALARM/WARN= 0/0
ALARM= 100% STOP= 40
=>
PF1=HELP 2=VRs 3=RETURN 4=SNI(HSCBS) 5=ADAPTERS 6=ROLL 8=CBPOOLS 9=NNTs
ENTER=REFRESH 10=UTIL 11=ALT-PFKS 12=REFRESH SUMMARY PA1=EXIT PA2=LOG

```

CCL offers new and enhanced security options for the traditional SNA workload

- **IP-based security for IP-partner authentication and data flow encryption can be extended to the IP-based CCL connectivity options:**
 - ▶ XOT and DLSw flows into CCL can be protected using standard IP Security (IPSec/VPN)
 - Between CCL and the partner XOT or DLSw router in the IP network.
 - ▶ INN or SNI connectivity to partner CCL NCPs over IP-TG can be protected using SSH tunnelling or standard IPSec Security (IPSec/VPN)
 - IP-TG between business partners require a single TCP connection between the two partner CCL NCPs with configurable listening TCP port numbers and IP addresses, which simplifies firewall configuration significantly

- **SNA Session Level Encryption (SLE) continues to be an SNA-based security option**
 - ▶ SLE can be used at an SNA session level through a CCL NCP as before

CCL performs better than an IBM 3745/46 ever did

- **Significantly higher transactional throughput than compared to a real IBM 3745/46 environment**
- **Consistent low response times even at high throughput rates**
- **From a CPU workload perspective, it is possible to consolidate up to five 70% utilized IBM 3745 31A NCPs onto a single System z9 IFL engine:**
 - ▶ In one Linux image with five CCL engines
 - ▶ In five Linux images, each with one CCL engine
- **Based on the better performance, it may be possible to consolidate multiple NCPs into fewer NCPs**
 - ▶ Keep in mind that an NCP remains limited to 16 Mbyte of memory, which is an architectural limitation in the IBM 3745 instruction set
 - ▶ Consolidating more NCPs into fewer NCPs may result in buffer shortages in the NCP
- **CCL can use today's high-speed LAN technologies**
 - ▶ Up to 10 Gbps for SNA traffic
 - ▶ No longer limited to the 16 Mbps speed of token-ring LANs

System Requirements and Reference Information

CCL requirement for System z hardware

➤ Processor support

- ▶ G5/G6, z800/z900, z890/z990, or System z9

➤ CP requirements (can be IFL engines on zSeries and System z9)

- ▶ Depends on workload and connectivity options

➤ OSA port requirement

- ▶ Copper-based ports for SNA LLC2 (LCS) - can be used on all hardware levels
- ▶ Fiber optic or copper ports for SNA LLC2 (QDIO layer-2) - z890, z990, System z9 only
- ▶ Fiber optic or copper ports for SNA over IP such as IP-TG, XOT, or DLSw (QDIO layer-3, QDIO layer-2, or LCS)
- ▶ OSN port for CDLC connectivity - System z9 only

➤ Memory requirements

- ▶ Memory per CCL engine: 20 MB
- ▶ Usual memory requirements for Linux on zSeries
 - Memory: 256 - 512 MB memory (depending on distribution, packages, and kernel level)

➤ DASD requirements

- ▶ DASD for CCL and InstallShield code = 65 MB
- ▶ DASD for Linux kernel source = 300 MB
- ▶ DASD for CCL traces, dumps, logs, NCP load modules = 80-100 MB per CCL Engine instance
- ▶ Usual DASD requirements for Linux on zSeries
 - Approximate DASD space equivalent to two 3390-3 DASD volumes
 - Use the Linux Logical Volume Manager (LVM) to group the volumes together

CCL requirements for Linux on System z

➤ **Minimum Linux requirements for CCL V1.2.1**

- ▶ SUSE LINUX Enterprise Server 8 for IBM zSeries and IBM S/390 (SLES8), kernel 2.4.21
 - Minimum level supported: Service Pack 4 (SLES8 + SP4)
 - **Note:** A Linux kernel level 2.4 does not support CDLC and QDIO layer 2 connectivity.
- ▶ SUSE LINUX Enterprise Server 9 for IBM zSeries and IBM S/390 (SLES9), kernel 2.6.5
 - Minimum level supported: Service Pack 1 (SLES9 + SP1)
 - Service Pack 3 (SLES9 + SP3) includes QDIO layer 2 and CDLC support (***The recommended level***)
- ▶ Red Hat Enterprise Linux AS 4 (RHEL4), kernel 2.6.9
 - Minimum level supported: Update 1 (RHEL4 + Update1)
 - Update 3 (RHEL4 + Update 3) includes QDIO layer 2 and CDLC support (***The recommended level***)
- ▶ Both 31-bit and 64-bit distributions are supported

➤ **Minimum Linux requirements for CCL V1.2.1 communication via QDIO layer 2 and CDLC**

- ▶ A Linux kernel level 2.6 is required
 - SUSE SLES 9 Service Pack 3 (SLES9 + SP3)
 - Red Hat AS 4 Update 3 (RHEL4 + Update 3)
- ▶ Processors:
 - For QDIO layer 2: IBM System z9 or IBM eServer zSeries z890, z990
 - For CDLC: IBM System z9

➤ **For availability of further distributions supporting CCL V1.2.1 functions and specific package requirements on top of available distributions refer to:**

- ▶ <http://www.ibm.com/software/network/ccl>

CCL is not a complete replacement for the IBM 3745/46 Communication Controller

CCL Functional Overview Matrix	CCL V1.2.1 supports	CCL V1.2.1 support of serial lines via an aggregation layer router	CCL V1.2.1 does not support
<p>Software</p>	<p>NCP (V7R5 and above) and compatible levels of NRF</p> <p>SSP, NTuneMON, NetView, and NPM continue to work as they have in the past</p> <p>NCP Packet Switching Interface (NPSI)</p>		<p>Other IBM 3745 software products: XI/NSF, EP, NTO, NSI, MERVA, and TPNS</p> <p>Functions provided by the IBM 3746 MAE or NNP (most of these functions can be migrated to CS Linux on System z)</p> <p>NCP-based IP routing (migrate to standard Linux-based IP routing)</p>
<p>Physical network interfaces</p>	<p>SNA LLC2 (LAN) access to OSA token-ring and Ethernet LAN</p> <p>NCP TIC2 or TIC3 LAN interfaces via OSA LCS or OSA QETH (QDIO layer-2)</p> <p>CDLC channel connectivity through shared OSA-E2 on System z9</p> <p>IP-TG for direct IP connectivity between two CCL NCPs</p> <p>XOT for x.25 connectivity</p> <p>DLSw for DLSw termination in Linux for System z</p>	<p>SDLC, Frame Relay, X.25 QLLC, and ISDN serial line interfaces are not supported directly by CCL, but are supported via an aggregation layer router</p> <p>X.25 circuits are not supported directly by CCL, but are via an aggregation layer router that uses the XOT protocol to transport the X.25 packets to/from NPSI running in CCL</p>	<p>BSC, ALC, Start/Stop</p>

Contact information for CCL

➤ CCL on the Web:

- ▶ CCL home page: <http://www.ibm.com/software/network/ccl>
- ▶ CCL news group: <news://news.software.ibm.com/ibm.software.linux.ccl>

➤ For more information, contact:

- ▶ EMEA: Peter Redman - Peter_Redman@uk.ibm.com
- ▶ North America: Erika Lewis - erika@us.ibm.com
- ▶ Latin America: Suvas Shah - suvas@us.ibm.com
- ▶ AP: Chuck Gardiner - cgardine@us.ibm.com



➤ For planning and installation services, contact:

- ▶ Heather Johnson in IBM SWG Application and Integration Middleware Software e-Server Services - hjd@us.ibm.com

➤ For technical assistance in the Americas, IBMers can submit a TechExpress through w3.ibm.com or a question through WWQ&A

For further technical assistance:

US:

- ▶ Access installation and technical support information via the WWQA database
 - IBMers can access via the WWQA database via QASearch on <http://w3.viewblue.ibm.com>
 - Customers can access installation and technical support information from IBMLink/ServiceLink.
- ▶ Please research questions through all available resources before submitting a question to the Q&A database.

EMEA

- ▶ Techline and local Field Technical Support Specialists provide technical pre-sales assistance. Additional technical support is available through worldwide Question & Answer (WWQA), QASearch function on ViewBlue or EHONE. For some brands/products, authoring of questions is only available via Techline.

Important CCL configuration samples

➤ Go to the CCL home page:

- ▶ <http://www.ibm.com/software/network/ccl>

IBM Software - Communication Controller for Linux on System z - Library - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

http://www-306.ibm.com/software/network/ccl/library/

Getting Started Latest Headlines WRAL.com - Weather Politiken.dk - det leven... Forside Jyllands-Posten washingtonpost.com - ...

Support

- Network Management for IBM Communication Controller for Linux on System z (CCL) Data Link Switching (DLSw)
- Anura Guruge on CCL V1.2
- Communication Controller replacement strategy
- IT In-Depth: IBM's Communications Controller for z/linux

Related software

- Communications Server
- z/OS Communications Server
- Network Control Program

Related hardware

- IBM System z
- IBM Communication Controllers
- Warranty info

Planning materials

- Capacity planning for SNI and boundary function
- CCL Frequently Asked Questions
- Supported features cross reference
- Supported Hardware
- Supported Kernels
- Installation and documentation notes

Brochures

- [Communication Controller for Linux on zSeries V1.2.1 Data Sheet \(170KB\)](#)
- [Communication Controller for Linux on zSeries V1.2 Data Sheet \(224KB\)](#)
- [Communication Controller for Linux on zSeries V1.1 Data Sheet \(271KB\)](#)

Redbooks

- IBM Communication Controller for Linux on System z V1.2.1 Implementation Guide
- IBM Communication Controller Migration Guide
- z/VM and Linux on zSeries: From LPAR to Virtual Servers in Two Days

Configuration samples

A collection of presentatons which provide sample configurations files to help you get started with CCL. For a list of the available samples, [click here.](#)

About IBM | Privacy | Contact

http://www.ibm.com/support/search.wss?tc=SSRRLB&rs=2192&dc=DB520&rank=8&q=cclconfigsampl

And click this link down here at the bottom of the page.

Configurations for a number of CCL connectivity samples

IBM - Search results - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

http://www-1.ibm.com/support/search.wss?tc=SSRRLB&rs=2192&rank=88&q=cclconfigsampl&dc=f

Getting Started Latest Headlines WRAL.com - Weather Politiken.dk - det leven... Forside Jyllands-Posten washingtonpost.com - ...

1 - 10 of 25 items found

		Modified date
[1]	Configuration sample for Ethernet INN using Cisco DLSw Sample configuration for ethernet connections using Communication Controller for Linux on zSeries (CCL). [More items like this found in Enterprise Connectivity] [This item's topic: Configuration]	2006-06-20
[2]	Configuration sample for Token Ring INN using Cisco DLSw Sample configuration for token ring connections using Communication Controller for Linux on zSeries (CCL). [More items like this found in Enterprise Connectivity] [This item's topic: Configuration]	2006-06-20
[3]	Configuration sample for Ethernet BNN using Cisco DLSw Sample configuration for ethernet connections using Communication Controller for Linux on zSeries (CCL). [More items like this found in Enterprise Connectivity] [This item's topic: Configuration]	2006-06-20
[4]	Configuration sample for SDLC INN Using Cisco DLSw Sample configuration for SDLC INN connections using Communication Controller for Linux on zSeries (CCL). [More items like this found in Enterprise Connectivity] [This item's topic: Configuration]	2006-06-20
[5]	Configuration sample for QLLC BNN using Cisco DLSw Sample configuration for QLLC connections using Communication Controller for Linux on zSeries (CCL). [More items like this found in Enterprise Connectivity] [This item's topic: Configuration]	2006-06-20
[6]	Configuration sample for Token Ring to Ethernet BNN using Cisco DLSw Sample configuration for converting Token Ring BNN to Ethernet using Communication Controller for Linux on zSeries (CCL). [More items like this found in Enterprise Connectivity] [This item's topic: Configuration]	2006-06-20
[7]	Configuration sample for SDLC BNN using Cisco DLSw	2006-03-13

Done

CCL release summary

➤ **CCL V1R1 (initial release) - March 2005**

- ▶ CCL NCP V7R5+, NRF, and MOSS console support
 - NCP SNI/INN support
 - NCP boundary function support
- ▶ OSA LCS SNA LLC2 connectivity - NCP TIC2 LINE addresses only
 - Copper based cabling
 - LAN speeds up to 1000 Megabit (1000BASE-T Ethernet)

➤ **CCL V1R1+ (APAR number LI70826) - August 2005**

- ▶ CCL emulator performance improvements

➤ **CCL V1R2 - November 2005**

- ▶ CCL support for NPSI
 - Connectivity to NPSI via XOT protocol handler from Eicon
- ▶ OSA QETH QDIO Layer-2 mode support
 - Fiber based cabling
 - LAN speeds up to 10 Gigabit
- ▶ CDLC connectivity on System z9 over shared OSA-E2 in OSA for NCP mode
- ▶ IP Transmission group for INN/SNI connectivity to partner CCL NCP
- ▶ Additional CCL emulator performance improvements

➤ **CCL V1.2.1 - May 2006**

- ▶ Imbedded DLSw support
- ▶ Native LAN support - TIC3 adapters and NCP TIC3 LINE addresses
- ▶ Local IP address control for IP-TG connections

For more information....



URL	Content
http://www.ibm.com/servers/eserver/zseries	IBM eServer zSeries Mainframe Servers
http://www.ibm.com/servers/eserver/zseries/networking	Networking: IBM zSeries Servers
http://www.ibm.com/servers/eserver/zseries/networking/technology.html	IBM Enterprise Servers: Networking Technologies
http://www.ibm.com/software/network/commserver	Communications Server product overview
http://www.ibm.com/software/network/commserver/zos/	z/OS Communications Server
http://www.ibm.com/software/network/commserver/z_lin/	Communications Server for Linux on zSeries
http://www.ibm.com/software/network/ccl	Communication Controller for Linux on zSeries
http://www.ibm.com/software/network/commserver/library	Communications Server products - white papers, product documentation, etc.
http://www.redbooks.ibm.com	ITSO Redbooks
http://www.ibm.com/software/network/commserver/support	Communications Server technical Support
http://www.ibm.com/support/techdocs/	Technical support documentation (techdocs, flashes, presentations, white papers, etc.)
http://www.rfc-editor.org/rfcsearch.html	Request For Comments (RFC)