



# An Introduction to the Communication Controller for Linux® on zSeries® Version 1 Release 1



CCL V1R1



## 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:

- ▶ AIX®
- ▶ AnyNet®
- ▶ AS/400®
- ▶ Candle®
- ▶ CICS®
- ▶ CICSplex®
- ▶ CICS/ESA®
- ▶ DB2®
- ▶ DB2 Connect™
- ▶ DPI®
- ▶ DRDA®
- ▶ e business (logo)®
- ▶ ESCON®
- ▶ eServer™
- ▶ ECKD™
- ▶ FFST™
- ▶ GDDM®
- ▶ GDPS®
- ▶ HiperSockets™
- ▶ IBM®
- ▶ Infoprint®
- ▶ IMS™
- ▶ IP PrintWay™
- ▶ iSeries™
- ▶ Language Environment®
- ▶ MQSeries®
- ▶ MVS™
- ▶ MVS/ESA™
- ▶ NetView®
- ▶ OS/2®
- ▶ OS/390®
- ▶ Parallel Sysplex®
- ▶ PrintWay™
- ▶ PR/SM™
- ▶ pSeries®
- ▶ RACF®
- ▶ Redbooks™
- ▶ Redbooks (logo)™
- ▶ S/390®
- ▶ System/390®
- ▶ ThinkPad®
- ▶ Tivoli®
- ▶ Tivoli (logo)®
- ▶ VM/ESA®
- ▶ VSE/ESA™
- ▶ VTAM®
- ▶ WebSphere®
- ▶ xSeries™
- ▶ z/Architecture™
- ▶ z/OS®
- ▶ z/VM®
- ▶ zSeries®

- ▶ 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.
- ▶ 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](http://www.ibm.com/legal/us) for further legal information.

## Today's Speaker



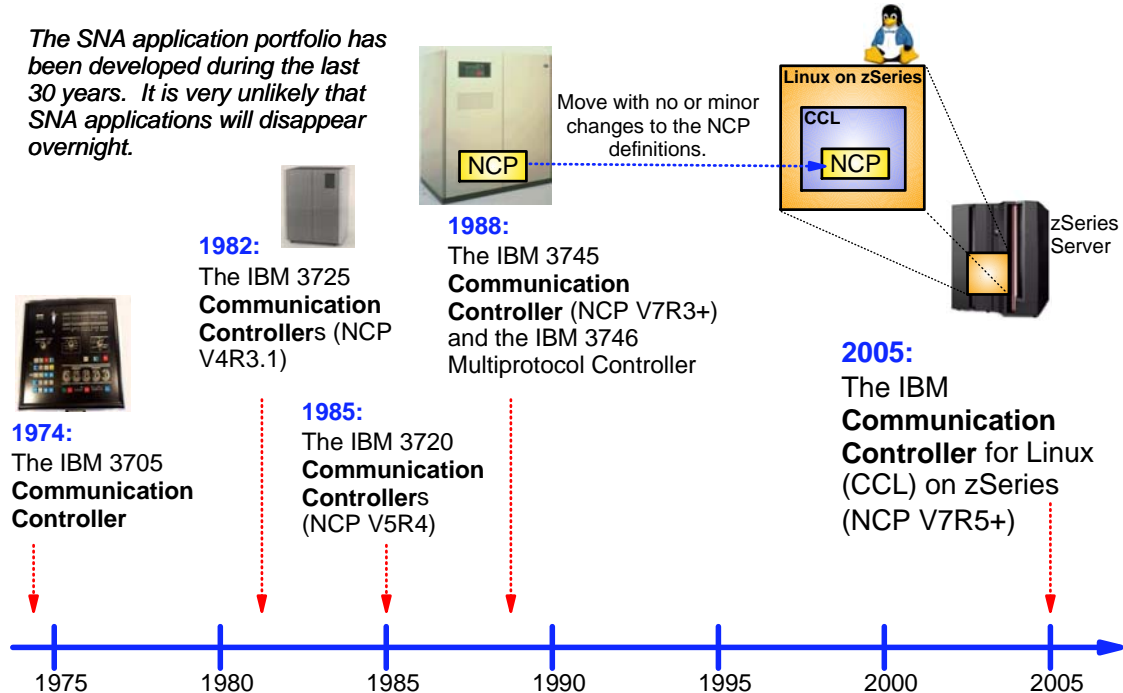
### > Alfred Christensen

- ▶ Almost 30 years of systems programming level experience on the IBM mainframe platform
- ▶ Currently a programming consultant in the Strategy and Design group within the Enterprise Networking Solutions area
- ▶ Primary focus areas are networking technologies on z/OS and integration of SNA and TCP/IP on Linux
- ▶ Author of numerous white papers, articles and IBM Redbooks
- ▶ Frequent presenter at SHARE, GSE and various IBM conferences

## IBM Communication Controllers - the foundation of SNA application access to the IBM mainframe since 1974



*The SNA application portfolio has been developed during the last 30 years. It is very unlikely that SNA applications will disappear overnight.*




## Do we still need a Network Control Program (NCP) in 2005?



- **The IBM 3745/46 hardware was withdrawn from marketing in September 2002.**
  - The hardware is currently still supported and serviced by IBM (plans to continue doing so to 2010)
  - The NCP software was not withdrawn and it is still currently supported and serviced by IBM.
- **The Network Control Program (NCP) runs on the IBM 3745/46 hardware and continues to be an integral part of a traditional SNA network infrastructure:**
  - Boundary functions for peripheral devices such as ATMs, terminal controllers, PC-based servers, etc.
  - SNA Network Interconnection (SNI) functions for SNA-based business partner communication.
- **Networking technologies to help remove dependency on an NCP have been made available, but implementing some of those technologies can be both time consuming and costly:**
  - Moving boundary functions to other platforms, such as VTAM, can be an administrative challenge.
  - SNA/IP integration products are abundant. They all require some amount of network infrastructure changes.
  - Upgrading to newer SNA architecture levels, such as Advanced Peer to Peer Networking (APPN), High Performance Routing (HPR), and eventually Enterprise Extender (EE), have many benefits but requires solid SNA skills, SNA network topology changes, and detailed coordination of network definitions across both an internal network and business partner networks.
- **Not all zSeries operating systems that today are using the services of an NCP have appropriate alternative technologies available:**
  - Only z/OS supports Enterprise Extender. VSE/ESA, z/VM, and TPF do not support EE technology.
  - SNI still requires at least one NCP.
- **SNA applications are still abundant and many of those SNA applications will most likely exist for years to come.**
  - Justifying the cost of rewriting SNA applications to IP-based applications just for the sake of removing NCP dependency is in many cases difficult.

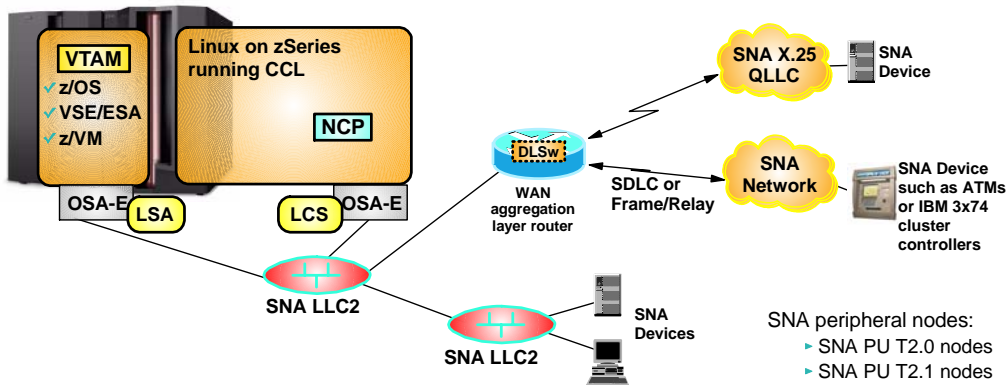
**YES, an NCP continues to be of value in many data center installations.**

## CCL Release 1 - Preserve SNI connectivity to business partners (SNI/INN traffic)



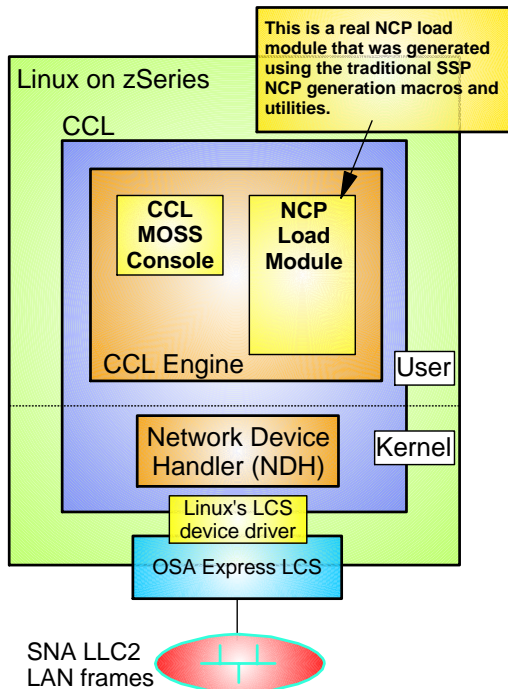
- **NCP SNI functions move to Linux on zSeries. Business partner may continue to use IBM 3745/46 technology or move to a CCL implementation also.**
  - SNA traffic leaves/enters the Communication Controller for Linux on zSeries as SNA network flows (SNA LLC2) over an OSA adapter operating in LCS mode.
  - VTAM sees the NCP as a LAN-attached remote NCP over its LSA OSA adapter.
  - SNA traffic can be tunneled (typically DLSw) over an IP network to the business partner's location.
  - An SDLC line from the business partner's IBM 3745/46 can be terminated in a local wide area network aggregation layer router (a router with WAN network interfaces).
  - Has no impact on existing SNI topology.

## CCL Release 1 - Preserve selected NCP boundary functions (BNN traffic)



- SNA wide area network links, such as SDLC, F/R, and SNA X.25 QLLC termination, moved from the IBM 3745/46 to a wide area network aggregation layer router that switches the SNA frames between the lines and the local SNA LAN
  - SDLC, Frame Relay, and SNA X.25 QLLC links are supported
    - The X.25 SNA support (QLLC) does not imply full NPSI support
  - NCP boundary function support includes standard availability functions such as SSCP takeover, support for duplicate MAC addressing, and XRF
  - NPA-LU, NtuneMON, and NRF are also supported
- Remote SNA LANs can be connected to the data center LAN by bridges or DLSw technology

## CCL structure and components



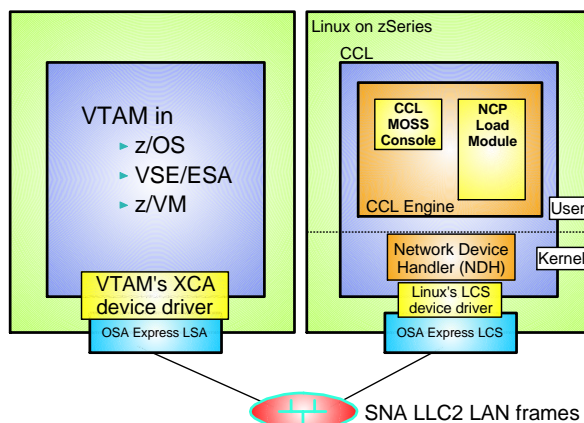
- CCL consists of both user-space and kernel-space functions:
  - **CCL engine** emulates an IBM 3745-31A with 16 MB memory supporting an NCP load module and a MOSS console interface.
  - The **MOSS console** is accessed through a standard Web browser.
  - **Network Device Handler (NDH)** is a kernel extension that acts as the interface between an OSA port operating in LCS mode and the NCP Token-Ring Interface (NTRI).
    - The only supported network interface from an NCP perspective is an SNA TIC2 interface.
    - The actual LAN to which the OSA port is connected may be either token-ring or IEEE802.3 Ethernet (NDH will transform between the frame formats).
- **NDH components**
  - NDH itself consists of two components:
    - A small source code isolation module that is built during installation of CCL
    - An object code only NDH module
  - Both are dynamically loaded into kernel space. No kernel rebuild/reboot is required.

## CCL and the MOSS console interface



Status	X71	X72	LAR	IAR	Level	C-Latch	Z-Latch
Running	000000	000000	13E38A	13E38E	2	1	0

## CCL and VTAM communication



➤ VTAM sees the CCL NCP as a LAN-attached remote NCP through its XCA (eXternal Communications Adapter) interface, such as an OSA LSA interface

➤ VTAM may either be the owning host or a data host to the CCL NCP

➤ For VTAM to be the owning host, a VTAM PTF will be made available for VTAM to activate and own a CCL NCP through an XCA network interface:

- New keyword on the XCA PU statement to allow VTAM to activate and own CCL NCP resources over an XCA interface
  - ALLOWACT=NO/YES

➤ 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

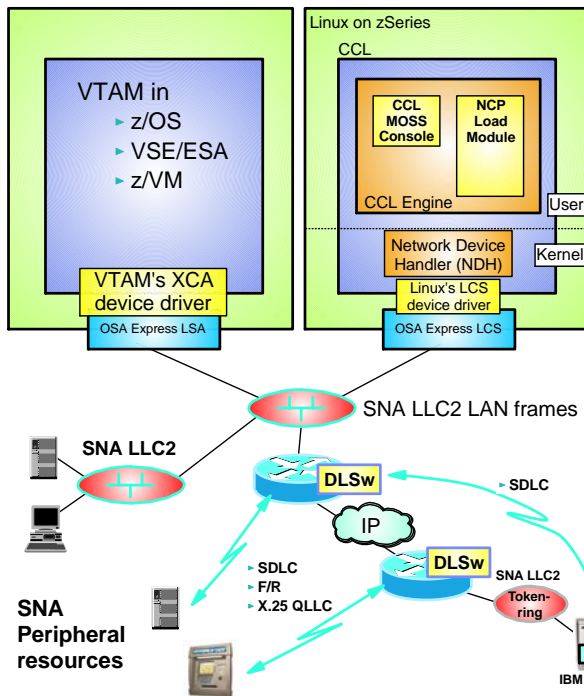
➤ None or minor changes to VTAM definitions and operations.

➤ In most cases no changes to NetView definitions and operations.

VTAM maintenance details:

- OS/390 and z/OS VTAM: APAR OA10425
- VSE/ESA VTAM: APAR DY46311
- z/VM VTAM: APAR VM63677

## CCL and downstream SNA connectivity



- Remote SNA lines need to be terminated in an aggregation layer router - a router with wide area network interface support such as a Cisco 3600 family router.
- The DLSw software in the router will internally switch the SNA PDUs between the wide area network link and LAN LLC2 frames:
  - ▶ Ethernet
  - ▶ Token-ring
- The NCP will see the remote resources as NTRI attached resources.
  - ▶ If migrating an NCP that today owns SDLC lines, the SDLC line definitions need to be changed to NTRI definitions.
- All NCP functions for PU Type 2.0 or PU Type 2.1 nodes are supported in the CCL NCP environment - including XRF sessions.

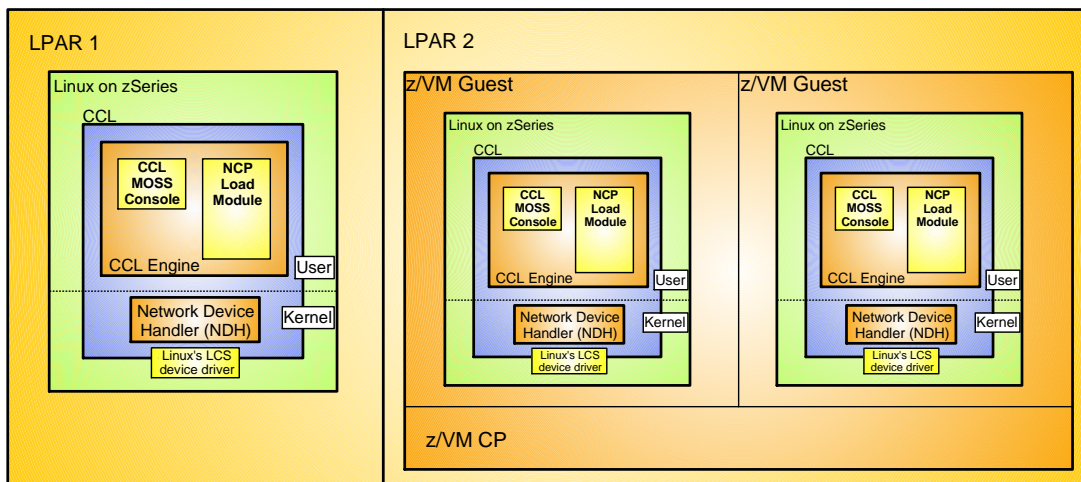
## Linux deployment model



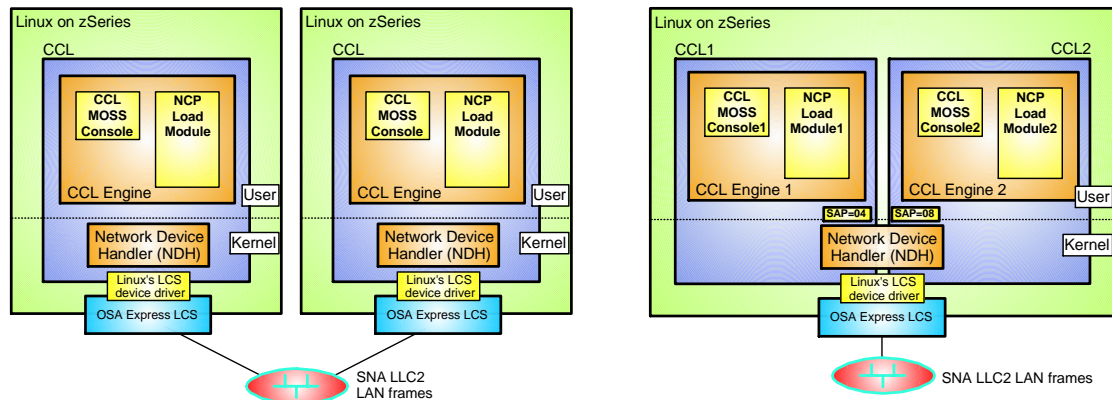
### ➤ Linux deployment:

- ▶ LPAR mode - one Linux image in an LPAR (no requirement for z/VM)
- ▶ As a z/VM guest

zSeries Server



## Deployment models for multiple NCPs



- One CCL NCP per Linux image
    - Each NCP operates completely independent of other NCPs
    - Each NCP may run in Linux images that are guests under z/VM or individual LPARs
  - Multiple CCLs per Linux image
    - One NCP per CCL
    - Each CCL has its own MOSS console
    - Multiple CCLs share one NDH instance
    - See OSA LCS port sharing rules later for details on using the same OSA LCS port by more NCPs
- ✓ No major difference in throughput
  - ✓ Two Linux images may require more overhead (DASD, Memory, CPU) than one
  - ✓ One Linux image may become a single point of failure

## CCL requirement for zSeries hardware



- **Processor support**
  - G5/G6, z800/z900, or z890/z990
- **CP requirements (can be IFL engines on zSeries)**
  - Depends on workload
    - SNI throughput 14% higher than a fully loaded IBM 3745-61A uses roughly one CP
    - SNA boundary function workload throughput 72% higher than a fully loaded IBM 3745-61A uses roughly one CP
- **OSA port requirement**
  - OSA-2 or OSA Express copper-based ports
    - IEEE802.3 Ethernet including 1000BaseT (10/100/1000 Mb)
    - Token-ring (4/16/100 Mb)
- **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: 50 MB
  - DASD for CCL traces, dumps, logs, NCP load modules: 80 - 100 MB per CCL engine
  - 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 zSeries



### > CCL runs under Linux on zSeries - initially the following distributions will be supported for CCL R1:

- ▶ SuSE
  - SuSE Linux Enterprise Server 8 for IBM zSeries and IBM S/390 (SLES8)
    - Recommended level is Service Pack 4 (SP4)
  - SuSE Linux Enterprise Server 9 for IBM zSeries and IBM S/390 (SLES9)
    - This is a Linux 2.6 kernel distribution
    - Recommended level is Service Pack 1 (SP1)
- ▶ Other distributions
  - We have tested with other Linux distributions and are working with distributors to incorporate an updated LCS device driver in their officially supported distributions.
  - At this time, Red Hat has not yet shipped a Red Hat distribution that includes the updated LCS device driver

### > Both 31-bit and 64-bit mode distributions are supported

### > There are specific package requirements to be met depending on the specific service level of the above distributions

## NCP support by CCL R1



- > **The CCL emulates an IBM 3745-31A with 16 MB memory**
  - ▶ Set MEMSIZE to 16MB on the BUILD statement to use all 16 MB
- > **The following NCP levels can be used to generate an NCP for the CCL:**
  - ▶ NCP V7R5
  - ▶ NCP V7R6
  - ▶ NCP V7R7
  - ▶ NCP V7R8
  - ▶ NCP V7R8.1
- > **SSP (System Support Program), NRF (Network Routing Facility), and NTuneMON at levels supported by the above NCP levels are also supported by CCL R1.**
- > **Only NCP Token Ring Interface (NTRI) is supported by an NCP running in the CCL**
  - ▶ Any device that can attach to an NCP over a TIC2 interface should be able to attach to an NCP running in the CCL environment
- > **The following native network attachments are specifically *not supported* by CCL:**
  - ▶ Channel resources
  - ▶ SDLC resources (use an aggregation layer router to bridge to LAN LLC2)
  - ▶ BSC resources
  - ▶ Start/Stop resources
  - ▶ Frame Relay resources (use an aggregation layer router to bridge to LAN LLC2)
  - ▶ EP resources in general
  - ▶ X.25 NPSI resources (for SNA X.25 QLLC: use an aggregation layer router to bridge to LAN LLC2)



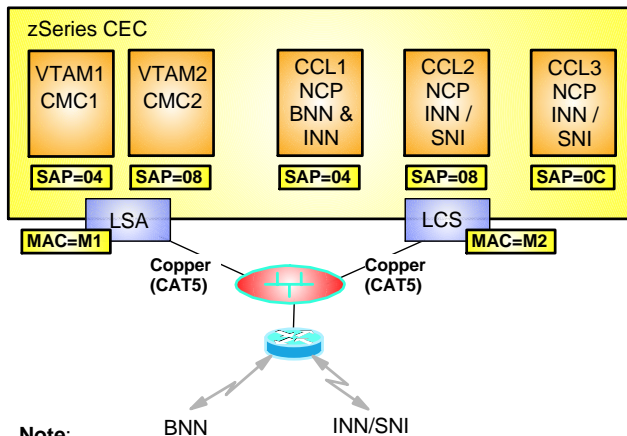
## CCL requirements for OSA ports



- CCL R1 exchanges SNA network flows with the network over a Linux LCS device driver interface only:
  - For NCP to VTAM communication (VTAM attached to shared LAN via an OSA LSA port)
  - For downstream communication where aggregation layer routers switch SNA PDUs to/from wide area network connections or over IP networks (DLSw)
- Only OSA copper-based interface ports can be configured as LCS ports - not fiber-based ports
- OSA/SF is needed for locally administered MAC addresses on OSA 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)

Processor type	Required MCL level	Ethernet (OSA Express FCs)	Token-Ring (OSA-2 and OSA Express FCs)
G5/G6		OSA Express FC 2340 Fast Ethernet (10/100 Mb) - 1 port/feature	OSA2 ENTR card FC 5201 (4/16 Mb) - 2 ports/feature (each port can also be configured as a 10 Mb Ethernet port)
z/800 or z/900	3.5	OSA Express FC 2366 Fast Ethernet (10/100 Mb) - 2 ports/feature	OSA Express FC 2367 (4/16/100 Mb) - 2 ports/feature
z/890 or z/990	5.50	OSA Express FC 1366 (upgraded) 1000BaseT (10/100/1000 Mb) - 2 ports/feature	OSA Express FC 2367 (4/16/100 Mb) - 2 ports/feature

## How many OSA copper ports?




**Note:**

1. The example shows all three CCLs in the same zSeries CEC. Separating them into separate zSeries CECs would mean each of them would need its own OSA LCS port.
2. For availability reasons, you need to consider doubling the number of ports for an implementation without single points of failure.
3. The two VTAMs can as always use SSCP takeover/giveback support to change ownership of the NCP resources.
4. For detailed OSA port usage planning information, *please see the CCL product documentation.*

- SNA LAN end-point is identified by a unique combination of MAC address and SAP number.
- An NCP BNN end-point must use local SAP 04.
- An NCP INN/SNI end-point can use a configurable local SAP for its subarea links.
- Multiple INN/SNI NCPs can share an OSA LCS port as long as they use unique local SAP numbers for their subarea links.
  - Use OSA/SF to add SAP number information to the OSA Addressing Table (OAT)
- Two or more BNN NCPs cannot share an OSA LCS port - they all require local SAP 04.
- VTAM (LSA) and CCL (LCS) cannot share an OSA port for communication between them.
- Two or more VTAMs can share an OSA LSA port as long as they use unique local SAP numbers and don't communicate with each other over that port.
- A Linux image can use the same LCS port for both SNA (CCL) and IP communication.
- Two CCLs that share an OSA port can use that shared port for INN traffic between them as long as they code UNIQUE=NO on their corresponding subarea PU definitions.

**In this case: Two ports**

## Disaster recovery site establishment



➤ For installations currently investing in DR sites, IBM 3745/46 redundancy poses some challenges.

➤ An alternative in many cases to installing spare IBM 3745/46 hardware will be to use Communication Controller for Linux running the DR NCP in the DR site.

➤ For LAN-attached connections, a switch to the DR site can be done using duplicate MAC addresses and layer-2 bridging of SNA flows, or using DLSw to redirect the traffic to the DR site.

➤ If physical serial lines are attached to the IBM 3745/46, they need as usual to be manually switched to the DR site where they then can be terminated in a DLSw router

➤ The CCL environment is able to reload and restart an NCP faster than a real IBM 3745/46, reducing recovery time of a failed NCP.

## IBM 3745/46 twin CCU (Communication Control Unit) configurations and the CCL environment



### ➤ Single CCU mode

- Only 1 CCU is installed in the controller
- Each CCL instance operates as a single CCU modeled after the IBM 3745 31A with 16MB memory

### ➤ Twin CCU in "dual mode"

- 2 CCUs installed in the controller
  - Channel and line adapters are dedicated to one CCU or the other
  - Bus switching between the 2 CCUs is not supported
- Similar functionality is achieved by running two CCL instances

### ➤ Twin CCU in "standby mode"

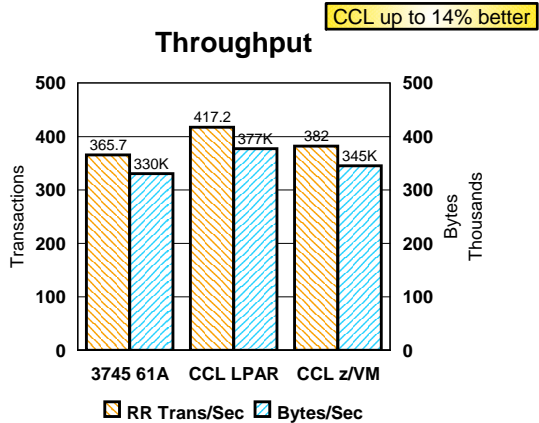
- 2 CCUs installed in the controller
  - All channel and line adapters are dedicated to only one CCU
  - The other CCU is down or idle ready to back up the other CCU
- Similar function is implemented within CCL itself
  - CCL will attempt to restart a CCL engine and load the same NCP load module as the failing CCL engine
  - In most cases, the automatic restart of CCL will happen more rapidly than with an IBM 3745/46 in twin CCU in "standby mode"

### ➤ Twin CCU in "backup mode"



- 2 CCUs installed in the controller
  - Channel and line adapters are dedicated to one CCU or the other
  - Bus switching between the 2 CCUs is supported for certain types of failures (power supply, CCU failures)
- Redundant hardware is provided by the zSeries platform and is available to Linux images running CCL



## SNI transactional workload - performance



For SNI transactional throughput equivalent to an IBM 3745 61A (or actually 14% better) Linux and CCL in these tests use one CP


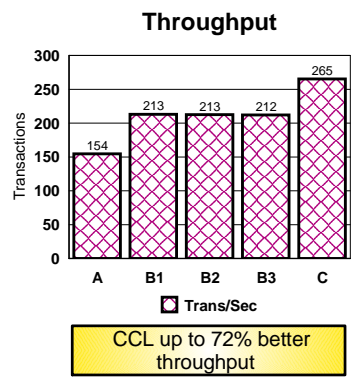



z/OS may see a slight CPU increase due to use of LSA instead of CDLC protocols when communicating with the NCP.

Workload details:

- 3270 transactions
- 100 bytes in / 800 bytes out
- z/OS systems running on a z900, CCL on a z990

## Boundary function transactional workload - performance



CCL results in about 2% less CPU on the Data host

For BNN transactional workload - Linux and CCL in these tests use one CP

- > **Config A:** CMC and data host attached via ESCON channels to an IBM 3745/46 (an IBM 3745-61A)
- > **Config B:** Boundary functions in VTAM CMC. Data host connected to CMC via:
  - Enterprise Extender over Fast Ethernet
  - SNA over a CTC running MPC channel protocols
  - SNA over a CTC - traditional VTAM to VTAM channel protocols
- > **Config C:** CMC and data host attached to CCL NCP over a Fast Ethernet

Workload details:

- 1 PU & 1 LU per device
- 3270 transactions
- 100 bytes in / 100 bytes out
- Think time 2.5 seconds
- z/OS systems running on a z900, CCL on a z990 under z/VM

## Terms and conditions



### > CCL pricing:

- ▶ Per processor pricing - \$40,000 per processor
- ▶ CCL available via Passport Advantage

### > NCP tier structure and pricing for the CCL environment:

- ▶ Each NCP running on a CCL requires a Tier 2 NCP license (\$620 per month)
  - Licensed to the zSeries serial number

### > Existing NCP terms and conditions also apply to the CCL environment:

- ▶ You need to maintain an NCP license for your IBM 3745/46 for as long as you continue to run an NCP on that hardware
- ▶ A new NCP license allows for a 60 day test period
- ▶ New NCP media shipment is based on NCP V7R8.1
- ▶ As usual you need an SSP license that matches the NCP level

#### Note:

1. Prices are current as of February 1, 2005, exclude applicable taxes, and are subject to change by IBM without notice.
2. Suggested retail price, dealer prices may vary.

## Contact information



### > CCL home page:

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

### > For more information, contact:

- ▶ EMEA: Peter Redman - [Peter\\_Redman@uk.ibm.com](mailto:Peter_Redman@uk.ibm.com)
- ▶ Americas: Erika Lewis - [erika@us.ibm.com](mailto:erika@us.ibm.com)
- ▶ AP: Chuck Gardiner - [cgardine@us.ibm.com](mailto:cgardine@us.ibm.com)



### > For planning and installation services, contact:

- ▶ April Singer in IBM Software Services for Websphere, Enterprise Transformation Services - [singeraf@us.ibm.com](mailto:singeraf@us.ibm.com)

#### 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](http://ibmlink.com/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.

## Summary



- CCL provides a migration path for SNA functions from the IBM 37xx hardware - allowing the NCP to be moved from a hardware communication controller, such as an IBM 3725, an IBM 3720, or an IBM 3745/46, to a Linux operating system environment on zSeries.
- This migration can be done with no or minimal impact to:
  - ▶ **The SNA network topology**
    - Retain subarea topology as it is today
    - Retain SNA path definitions as they are today
    - Retain SNI topology and connectivity as it is today
  - ▶ **VTAM definitions**
    - One new keyword on the XCA PU to allow VTAM to activate and own an NCP over a LAN interface
  - ▶ **Operation**
    - NCP generation and maintenance process is unchanged
    - Load and activation of the NCP might be slightly different from the way an installation has chosen to do it today, but the procedures are not new - they also work with an NCP that runs in an IBM 3745/46 environment
    - VTAM commands to manage the NCP and the resources owned by the NCP are unchanged
  - ▶ **Management of the NCP**
    - Same management interfaces as today.
    - NetView for z/OS, NPM, and NTuneMon can be used as today.
- If this is the first time an installation begins to use Linux on zSeries for production workload, some amount of management processes need to be established for Linux on zSeries.

And all of this is totally transparent to our SNA applications. No SNA application rewrite is needed to take advantage of this technology.



Thank you for joining us here today!

## For More Information



- Stay up to date on events with CCR2™, a monthly publication dedicated to zSeries systems management:
  - ▶ [www.ibm.com/software/tivoli/features/ccr2/info.html](http://www.ibm.com/software/tivoli/features/ccr2/info.html)
  - ▶ Select Subscribe at left
  
- Stay tuned for the next Tivoli zSeries webcast:
  - ▶ Check CCR2, The Mainstream or the zSeries Software Web page