z/JOURNAL

Alternatives for Replacing IBM's 37x5 for z/OS Users: A Comparison—CCL or Enterprise Extender?

BY GORDON WEBBER

Despite IBM's current policy of support until 2010, several questions concerning the viability of the 3745 (e.g., expertise, spare parts, etc.) will force the issue before then.

The question for the 3745 user today is, "What's the best strategy for replacing the 3745 and its functionality?" There are many factors involved in making that decision, such as where the solution should reside and which solution offers the best long-term progression.

IBM has two offerings for z/OS users and this article compares them:

- CCL for Linux
- Enterprise Extender.

Neither solution provides total support for *all* the physical connections and protocols the 3745 supports, but both solutions provide the most-used functions—TN3270 and SNI (System Network Interconnect).

For users of the old protocols—BSC, S/S, NPSI (except QLLC) XI, etc.—there's little alternative but to replace the service. The user base is so small there's no commercial incentive to produce an alternative solution. For some of these services, such as BSC, protocol converters may provide a short-term answer.

IBM has been coaxing users away from those old protocols for some time. Most have done so or are at least planning to. Certainly, most users of rare protocols (such as airline protocols) have already sought alternatives.

Even restricting comment to just the supported protocols, several factors could be potentially involved in making a choice, meaning that it's simply impractical to cover all possible permutations. Instead, we'll try to look at some of the practical aspects of how each of these solutions can be established and what's involved in such an implementation.

Communications Controller for Linux (CCL)

CCL is what IBM describes as the logical solution. The premise that it's neither cost-effective nor practical to redevelop the SNA applications is entirely sensible and applicable to any solution.

CCL involves taking the existing Network Control Program (NCP) and, having removed the Front-End Processor (FEP), running the same code in a Linux Logical Partition (LPAR) or Linux guest under z/VM in a shell (emulator).

The early communications controllers were intended to offload the processing workload, especially the line-driving to another CPU. That CPU (CCU in FEP terms) was located in the channel-attached FEP.

For the purposes of this discussion, we can ignore its predecessors and say that the 3725 was really the first mainstay of the FEP market; it revolutionized the hardware and performance of these boxes, becoming the acme of the NCP-only world for many years.

What all 37x5 boxes have in common, though, is their unique bit-oriented architecture, single-purpose design, and lack of storage! The programming language of these systems (bit-oriented assembler) is quite distinct from the 360/370/390 assembler of the mainframes to which these boxes were connected, having been derived from the need to integrate with the hardware, save space, and efficiently manage this bit-oriented, line-driving environment.

This is important when you consider CCL isn't a straight port; it's only the code itself that's been ported. The line functions (a large part of the work the 37x5 performed) don't follow this solution to the mainframe and must be implemented on a completely different outboard platform (an aggregate router) with ensuing costs.

What CCL has done is to separate the operating system from its closely related hardware, and then create a shell in which to replicate this environment so the code can be executed unchanged. All this effort now occurs on the same mainframe that spawned the FEP because of its need to offload the work.

IBM's performance figures show a linear progression for CCL—close to that of an average 3745. Remember, though, this new environment doesn't include the line-driving to the same degree as the self-contained 3745. Also, the workload on the router that now terminates the physical lines must be accounted for.

On the face of it, this solution seems

ideal, since it requires no additional knowledge or technology (for the network) and should be a cost-effective solution. However, this won't be the case for most users.

Bearing in mind the 3745's self-contained environment, let's break down the construction of CCL:

VTAM: VTAM has to own an NCP to control its services. It's from this position (or possibly via NetView) that commands are sent to the FEP to load NCP modules, start and stop lines, take dumps, etc. While NCP commands remain the same, none of the physical connection hardware can be controlled from here, since the physical connections must be transferred to the aggregate router(s). Also, there are some small (depending upon the current network topology) changes required to actually connect VTAM to the CCL/NCP (e.g., an XCA [eXternal Communications Adapter] major node for the LSA [Local Security Access] connection to the 3745; the VTAM-CCL communication now occurs via the OSA [Open System Adapter] internal network or a Cisco CIP).

CCL LPAR: CCL can run in several combinations. Cost is a factor here because the CCL is licensed by CP, the general purpose central processor in the zSeries. If two CCL engines are required, there's a choice of either running both on the same CP (single charge; less performance) or using a different CP per CCL (double the charge).

The issue here is likely to be capacity; a well-utilized 3745 could equate to using more than half a CP, so the single CP option may not be possible.

IBM suggests that using a general CP wouldn't be cost-effective (see IBM CCL FAQ document CCLV1R1_FAQs 2005 [FQ109582]). Extra CPs mean extra MIPS and extra license costs. This solution should run under Integrated Facility for Linux (IFL), which provides Linux-only CPs at reduced MIPS cost. Those CPs don't count toward the MIPS-count of the processor. Of course, there's a dollar cost associated with the IFL.

NCP: CCL will definitely require changes to the NCP. The amount of change depends on the current and desired topology. The CCL communicates both to VTAM and the lines using the OSA. CCL uses NCP Token Ring Interconnect/Interface (NTRI) as its internal protocol, but either NTRI or Ethernet to communicate with the outside world, a conversion being provided by its Network Device Handler (NDH) component at the boundary of the CCL engine.

All direct line connections must be redefined as NTRI; this is how OSA handles locally attached SNA devices. NTRI is an end-of-life technology.

Since the local connection (NTRI) is defined using LAN Channel Station (LCS), it can't use Queued Direct Input/Output (QDIO) mode on OSA Express. It seems likely that the existing line configurations in the NCP will remain intact and the CCL/NDH provides some form of mapping to the NTRI VMACs, which are then used in the local DLSw at the aggregate router to locate the real physical link.

This leads to the question: How valid is the NCP display information regarding those physical connections? Since NDH answers such requests in a CCL/NCP environment, NDH will have to interrogate the router to get real status.

No performance information is available regarding the processor (what was the CCU) itself or the load on the Token Ring. However, some statistics of byte rate and packet counts are stored in a Linux file.

On a real 3745, the NCP license is based on the usage tier of the box. This is determined by the hardware configuration. With CCL, however, this will always be regarded as a usage tier 2 license, at least for the time being.

OSA: The whole premise of CCL depends on copper OSA connections. The communications between CCL and both VTAM and the physical circuits goes via the OSA (or possibly a CIP router for the VTAM connection) and this environment can't cope with Fibre or QDIO, so a copper card is needed.

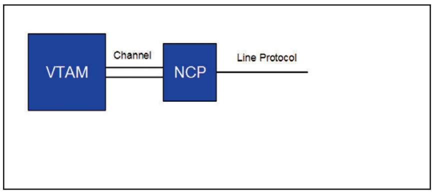

Figure 1: 3745 Connectivity

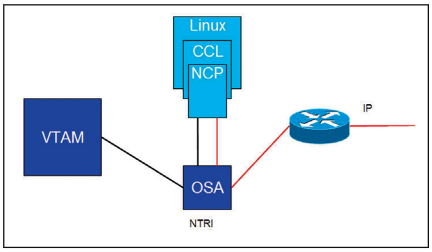

Figure 2: CCL Connectivity

You must configure this with OSA Support Facility (OSA/SF), which may not be currently in use and may be a new addition to your NCP administrator's skillset.

Apart from the need for copper, there are other restrictions that govern what combinations of devices can attach to the same OSA card, so additional OSA ports may be required.

For true resilience, you may need to add OSA ports.

Aggregate routers: Aggregation refers to the convergence of Layers 1 (physical), 2 (Link), and 3 (Network) traffic. These routers typically provide access via MAC or IP addressing. It's these routers that will now have to terminate the physical lines that used to be attached to the 3745. This may (and almost certainly will) incur cost in providing the correct interface cards in the routers and may even require dedicated routers for viewing the possible traffic loads.

A fully loaded 3745 can handle about 1,100 lines. How many of these must be

- Service (upgrades or PTFs) may be required to support the required environment.
- Some configuration changes to VTAM will be required.
- Some configuration changes to NCP will be required.
- The architecture will change; this isn't a direct port!
- Additional OSA ports may be required for the immediate solution or for resilience.
- Additional router hardware may be required, either routers themselves, or at least interface cards for the physical link attachments.
- The knowledge required to operate this system is greater than that previously required for an FEP/NCP combination, and those skills may not already be available.
- Further CPs may need to be purchased (general CP or IFL) at significant cost.
- The license costs could be high; more than one CCL/CP license may be required.
- Additional CPs could affect the MIP rating of the system and the license costs of all other products.
- Throughput could be affected as CCL brings the FEP workload back to the mainframe.
- Control is more difficult as the solution is now spread across multiple disciplines.
- Troubleshooting is also now spread across several diverse segments of the network.

Figure 3: CCL Summary

terminated at a router and how many routers will this take?

A redundant router will probably be required to provide sufficient resilience; otherwise, this will become a single point of failure. Even with redundancy, sessions may not be recovered if one router fails because session details aren't understood at this level. NCP links are often logically defined as Multi-Link Transmission Groups (MLTGs) to provide resilience against the failure of a single link in the group. To provide a similar facility for CCL, redundant DLSw equipment may be required.

These routers may require a separate monitor to keep track of the workload, throughput and define and manage the new interfaces.

These routers may bring another new skillset to the network, but even if the basic router operating system isn't new, it's likely the interfaces will be.

Tools

The topology has evolved from what's shown in Figure 1 to Figure 2. This means the skills and tools required to troubleshoot a connection include VTAM, NCP, OMVS, Linux, CCL, OSA, router technology and configuration, router diagnostics, SNA, and IP. These tools are no longer located in one place.

The Network

One of the main features that endeared SNI to its users is the security of the connection. NCPs connected between disparate networks (SNI) are always connected by private wires and the way that VTAM/SNA controls the connectivity across this link means that only designated users can access designated resources. This is a highly secure environment.

As soon as the NCP is moved onto the mainframe, that direct connection to the line is removed and a risk is introduced in the path through the aggregate router. The degree of the risk will be proportional to the distance to the router from the CCL.

DLSw in general doesn't take account of SNA class of service (used to prioritize SNA traffic, especially for interactive sessions and management. How local DLSw affects this isn't clear, but you should investigate this with the vendor of the router you select.

It may be tempting, once the NCP is connected to the aggregate routers, to convert away from the expensive SDLC, F/Relay or X.25 circuits and use the IP network or perhaps even the Internet. However, as soon as this happens, security becomes an issue as the data is all clear text.

Figure 3 provides a summary of key points about CCL.

Enterprise Extender

IBM, by its own admission, describes Enterprise Extender (EE) as the preferred solution for the most common functions of SNA supported from the NCP. EE provides a UDP encapsulation service that allows use of the IP infrastructure to carry the SNA frames to their destination (see Figure 4). EE has a far cleaner path than CCL, has been around for many years, is more efficient and is easier to support. It will provide connectivity for TN3270 and APPN (LU6.2, HPR, DLUr/DLUs), and covers the same functional area as CCL.

However, to implement EE, APPN/HPR is required, and the NCP functions will need to be disassembled and re-applied. This doesn't, of course, apply to the applications, which remain blissfully unaware of any transition from NCP to EE.

EE doesn't require the connections and interfaces to be migrated since it will replace the entire network path by the use of the IP infrastructure. EE does have the unique advantage of being able to maintain the SNA class of service across the IP network by use of the ToS header bits and the mapping of SNA priority levels to different (IANA registered) ports.

Let's look at the outline process required to implement EE:

APPN: APPN is often seen as the biggest hurdle to overcome to implement EE. For a user with a sub-area network, the conversion to APPN may seem daunting. However, APPN is no longer as problematic as it once was. Provided the correct groundwork is prepared (particularly regarding routing and ensuring a solid naming convention), it's not such a risk. If a user is already APPN-enabled, or is planning to convert anyway (when the conversion of the NCP circuits could be incorporated into the migration), then EE is surprisingly simple to activate.

If the user doesn't wish to convert, then it may be possible to set up a new, single-APPN network node to handle the EE connections (similar to the way the FEP offloads the communications workload) and run that node as an Inter-Connection Node (ICN), leaving the sub-area network almost intact.

The functionality that provides EE already comes with z/OS operating systems. VTAM is given an IP address via a VIPA, which gives the TCP/IP stack a path by which it can send traffic to VTAM. A device in the TCP/IP stack is used to drive the connection and VTAM XCA and SWNET major nodes need to be created to define the path to the remote system.

APPN/HPR provides the routing and recovery management at the SNA level, while the TCP/IP stack provides the link to any attached IP network.

- APPN will create the biggest work effort in moving to EE if the user isn't already APPN-enabled.
- EE is simple to install, robust, and reliable.
- EE is IBM's preferred solution and one that will continue to be supported and developed.
- Implementation is simple, with little or no disruption to current operations.
- No additional licenses or hardware should be required.
- No additional skills are required.

Figure 5: Enterprise Extender Summary

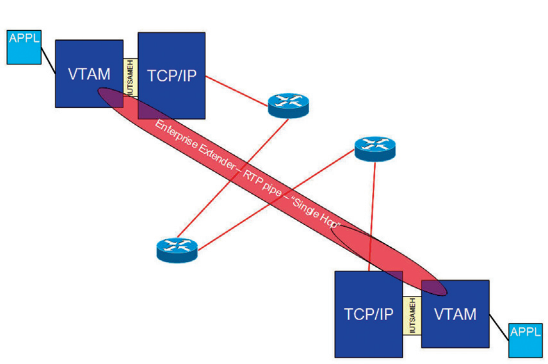

Figure 4: Enterprise Extender Connectivity

If APPN is enabled, the affected areas required to define this are:

- **VTAM:** Changes to ATCSTR00 and the addition of an XCA and SWNET major node
- **TCP/IP:** The definition of the device and VIPA in the TCP/IP parameters
- **CONSOLE/VTAM commands:** Vary commands activate and deactivate the link.

The Network

No changes are required to the IP network, although security should be considered as EE provides an unchecked path between each end (unlike SNI), allowing any user access to any resource. This can be achieved in several ways, including VTAM session-level encryption, IPSec VPNs, or a proprietary product that provides SSL authentication and encryption for EE links.

The EE solution may also use the Internet with no further consideration or change to the host environment (other than security). Figure 5 provides a summary of key points about EE.

Conclusion

CCL is a useful tool, especially for those users (such as VSE) who have no alternative. However, it should be regarded as a migration aid. While the original intent of CCL is valid and meaningful, users must accept the complexity and cost of CCL. For users who can move to EE, it makes more sense to do so now, as this would seem to be the ultimate long-term target.

For a z/OS site with a single link, spare CPU capacity, existing implementations of z/Linux and spare router capacity, CCL represents a sensible migration path. However, it's questionable whether the "effort" required to go EE is anything like the disruption or cost of implementing CCL for most z/OS shops.

As both EE and CCL are IBM products, it seems fitting to let IBM have the last word, with another extract from their CCL FAQ document:

Q: *Compare CCL with Enterprise Extender. Which alternative is a better way to migrate from SNA to an IP network?*

A: *EE is the better technology! EE is based on the latest SNA architectures, APPN and HPR, and supports IP end-to-end, all the dynamics of APPN, and the non-disruptive path switch of HPR. If both you and your business partner are at an APPN/HPR-capable level and agree to use EE between you, then this is the best and recommended way to go. CCL is for those who can't or won't migrate to the newer, better SNA architecture levels.* Z

About the Author

Gordon Webber is the product evangelist for William Data Systems. He has worked in the IT industry for 35 years (including time spent at IBM), covering roles in application and systems programming, product support and training, and has worked specifically in the network arena for the last 16 years. He has worked with mainframe systems, open systems, and security systems.
e-Mail: Gordon.webber@willdata.com
Website: www.willdata.com