FEP-4600 Communication Controller
A white Paper

FEP-4600 Communication Controller

Overview

Since 2000 the data center world began moving away from confining, proprietary mainframe solutions and towards non-proprietary, Open System solutions. Everyone will agree that change in technology is inevitable and in the long run is positive for all. However, change affects some differently than others, especially those dependent older mission critical protocols such as BSC, SNI and NPSI; this is our market.

The FEP-4600 Communication Controller focuses on the IBM® data center as a functional and very cost effective replacement for 3745/46 Communications Processors, more commonly known as front-end processors (FEP), that provide the mainframe user with the ability to perform high speed SNI, Bisync and LAN communications without the need for the traditional 3745/46 FEP and its associated costly software products.

Connecting directly to IBM and IBM-compatible mainframes via a direct channel connection, the FEP-4600 provides full communications between the mainframe and the remote network, which gives the customer SNA, Bisync, and PU4 communications while eliminating the need for costly licensed software from IBM. The robust architecture of the FEP-4600 allows units to be stacked using a single channel interface, thus providing an expandable system. Changes and reconfigurations can be performed while the equipment is fully operational so there is no need to bring the entire network down for maintenance.

The FEP-4600 is designed to specifically replace existing equipment, its target market is easily defined as any Data Center with a 3745/46 FEP utilizing SNI, BSC, NPSI and Token Ring

Most big companies still quietly operate FEPs; it is estimated that some 20,000 are still in operation. Whenever there's a merger, acquisition or a new disaster recovery plan, there's a FEP involved. No one likes to talk about it, but FEPs are still handling some important network activity. A good part of the worldwide financial industry continues to rely upon SNA and FEPs. Security and reliability still matter.

Paper Organization

The focus of this paper is to provide the reader insight into why the Visara FEP-4600 Communications Controller is the only viable replacement for the traditional IBM 3745/46 Front-end Processor and is organized as follows:

FEP-4600 Competitors; who are they, their product(s) strengths and weaknesses, how they compare to the FEP-4600. There are several solutions that offer z/OS users the opportunity for changing out their attached front-end processors. This Guide explains ands compare the following FEP replacement products.

- IBM's Communication Controller for Linux (CCL)
- IBM's Communication Server (CSL)/Enterprise Extender (EE)
- IBM used/Existing Front End Processors (FEP)
- Cisco Channel-attached Router 7200 series ECP4, 7500 series CIP (now in EOL)

Traditional Front-end Processor

History

IBM is the major competitor in the corporate data network market. Traditionally, the connection from the mainframe to the external data network is via the Front-end Processor (FEP). Historically, the first FEP was a 2701 that handled bisynchronous (Bisync) communications lines. With the advent of SNA communications, IBM introduced the 3705, which was superseded by the 3725, the 3720 and finally the 3745. These devices are also physically very large and control many communication lines. To provide configuration and control of these communication lines, the FEP runs several software programs. These are Advanced Communications Function/Network Control Program (ACF/NCP) typically known as just NCP and System Services Program (ACF/SSP). Depending on the FEP model and operating system, NCP and SSP are licensed at approximately $1,000 to $3,000 per month. If it is necessary to provide Bisync communications, then another software product, Partitioned Emulation Program (PEP), is required, which adds $350 per month to the operating costs. In large data centers that have multiple FEPs, each FEP has its own software license. NCP, besides being costly, is also very cumbersome and difficult to use. In fact, many large companies have employees dedicated to configuring and maintaining NCP, while smaller companies are forced to hire consultants to maintain their networks.

Life Expectancy

In 1999 IBM stated in a promotional document supporting the 3745/46/NCP that:

- Migration to all IP will take time and without the 3745/46 there is no easy alternative for SNI, BiSync and NPSI (X.25).
- The Internet's suitability for mission critical quality and security was only 58%.

Even though the above statements were made 7 years ago the pure fact is that, as then, there are still some 20,000 37X5 FEP's in use today is testament to their accuracy.

In 2002 IBM made the decision to end-of-life (EOL) the 3745/46 family of Front End Processors. As a consequence some firms found themselves caught in a technology void. Many where not, for one reason or another, positioned to make the change to either a non-proprietary or Open Systems solution. Firms most affected by the 3745/46 EOL were those running business critical applications connected via:

- SNI (SNA Network Interconnection)
- BSC (Bi-Synchronous Communications)
- NPSI (X.25)
- Token Ring

IBM has promised maintenance support at least through 2010, rumors suggest it may continue through 2015. IBM's support for NCP remains superb - it has to be, considering that NCP is not going away for many years. Hence, IBM's Communications Controller for Linux (CCL), an interim FEP software emulator which continues to rely upon NCP.

Most FEP users have already moved the all their easy stuff to their IP networks utilizing a Cisco CIP or CPA, which is another EOL story in itself (see Cisco below). What users found out then and is still true today, secure SNI using real FEPs is preferred. Other not so easily migrated applications such as EP, NPSI and SP also require a real FEP. So as long as spare parts and qualified service people are available and an IT budget willing support high software licenses; like the little pink rabbit the FEPs just keep on working!

Stratergy

If it wasn't for those that continue to require mission critical legacy protocols such as SNA/SNI, Bisync and NPSI; the phrase FEP Strategy would be meaning less. There wouldn't be any FEPs. So the phrase should read FEP Necessity. The only real strategy in is how to keep the cost of my SNA operation be reduced to a minimum. The smart way to do this is to consolidate FEP workload on to a modern FEP replacement platform with a short ROI and that does not require costly software licenses and features optimum reliability and serviceability characteristics. This is the easiest and lowest cost approach for legacy SNA/SNI solutions.

In the face of IBM's current policy to support 3745/46 FEPs through 2010 several questions concerning the viability of the 3745/46 exist. The biggest concern of the remaining Front-end Processor customers have confronted since the EOL announcement, was that a shortage of spare parts or personnel with true 3745/46 maintenance skills could leave them totally incapacitated the next time a FEP crashes, with no means of recovery. Many customers have succumbed to the belief that; "the next time your 3745 crashes could be its last".

The real question for the remaining 3745/46 users today is, "What's the best strategy for replacing the functionality of my Front-end Processor?" 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. Yes, of course, they can stay with their current hardware. But, neither IBM's nor a Cisco solution can provide total support for the physical connections and protocols. The only solution that offers the complete and total migration of a customer's FEP is the Visara FEP-4600 Communications Controller.

Excerpted from IBM's "Communications Controller Migration Guide"

> *Leaving things as they are (at least for now): Continuing to run an IT service without change can save your organization the cost and potential disruptions implicit in a migration. One situation in which such a strategy may make sense is when the use of the IT service in question is expected to decline substantially over time (perhaps due to business process reengineering efforts). This strategy, however, has the risk that at some point in time the hardware and software foundation of your configuration will become unsupported. Running on an unsupported configuration means that recovery from an outage may be difficult or impossible. You should make sure that your use of the configuration will end before support for its underlying products is terminated, or that loss of the service will not severely affect your business operation.*

While the option of leaving things as they are, above, suggests waiting for the more strategic projects to be complete, it might also be appropriate to try and expedite such projects by adding to their business case the potential savings from optimization and consolidation of your communications controller environment.

FEP-4600 Competition

IBM: Communications Controller for Linux (CCL)

CCL is a software only product that runs in a Linux partition on the mainframe. It requires and uses NCP, the same licensed product that runs on the 3745/46 FEPs. In addition, if X.25 is required, there is an additional IBM licensed software product required: NPSI. But wait, it gets even better. If you are running X.25, in addition to NPSI, you must also obtain a third party software product from Eicon that takes the output from NPSI and transports it to the external router that is required to drive the serial interface. The CCL solution is an expensive solution at best and very expensive solution at worst. The CCL is however feature rich, and provides real solid solutions.

Strengths:

- No traditional Front End Processor
- It enables SNA customers to maintain much of the unique, value-added networking functionality provided by ACF/NCP
- By removing the Front End Processor CCL provides NCP with a new lease on life
- CCL Token Ring implementation is very straight forward for older CPUs that support Token Ring OSAs (Open Systems Adapter).

Weakness:

- Does not support or address EP/BSC, Start-Stop, ALC.
- Requires two or more OSA's. OSA Support Facility (OSA/SF) must be configured, which may not be in use presently and could require a new NCP skill set.
- Newer IBM CPUs (System z9 and zSeries) do not support Token Ring OSAs, only Ethernet OSAs. To implement a Token Ring solution on newer mainframes would require enlisting a router to provide the Ethernet to Token Ring conversion.
- Ficon XOT Software Adapter – 3rd party software required in addition to NPSI for X.25 connections only, runs in same Linux partition
- External Routers required for external serial connections – most do not have large concentration of serial lines

Costs/considerations

- $40,000 one time cost associated with CCL (includes first years maintenance fee) for each CPU it is installed on $8000 annual maintenance fee for each CPU it is installed on Tier 2 NCP License fee (estimated at $630/mo) for each instance of CCL that is being run Tier 2 NPSI License fee (estimated to be higher than NCP license) for each instance of NPSI being run (only required for X.25 connections, runs in same Linux partition)
- SSP Software is still required to support NCP (estimated at $350/mo)
- Can consume large quantities of CPU resources. CCL V1R2 when configured to run flat out (SNI connection) consumes 55% of a z/990 CPU.
- Upgrading or adding to current CPU resources to run the CCL Linux partition.

CCL SNI Implementation:

The IBM recommended solution differs depending on what you are talking to in your partner's network. If the platform in the adjacent network is running a real 3745/46 FEP, then the CCL implementation would be to use DLSw (Data Link Switching - SNA over IP) directly from CCL through the OSA to an external router. The router would then convert the DLSw protocol back to SDLC to connect to the remote 3745/46 FEP. With this implementation there would be no multiple line transmission groups. The implementation should be comparable in performance to ours (limited to the speed of the SDLC connection).

If the platform in the adjacent network is also running CCL, then the recommended implementation would be to use IP Transmission Groups between the two CPUs. This would be implemented by connecting the OSA to a router/WAN connection to the remote site where the implementation would be the same. The IP Transmission Groups allows for a very high throughput of SNI traffic between the two networks, greater than what we can achieve through the FEP-4600 SDLC interface.

CCL Pricing Considerations

The US list price for CCL is $40,000 per System z9, zSeries, G5, or G6 processor. This means the price is dependent on the number of CPs running CCL, rather than the number of copies of CCL. For example, if you have one CCL running in two CPs, this would be $80,000. On the other hand, if you have one CP running two CCLs, this would be $40,000. Your implementation of CCL will depend on your Linux installation and configuration, your performance requirements, and so on.

The NCP is priced per NCP image as a 3745-31A running at Usage Tier 2. Each instance of CCL runs a unique NCP instance and therefore requires a separate NCP license. The CCL price includes the first year's maintenance. After year 1, the US list price for annual maintenance is $8,000 per System z9, zSeries, G5, or G6 processor. CCL's Yearly Cost of $40,000 is a one-time cost, including the first year's maintenance charge.

CCL Security

One of the main Front-End Processor features that locked SNI to its users is the security of the connection. NCPs connected between different 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 using CCL, that direct connection to the line is removed and a risk is introduced in the path through the aggregate router. The amount of the risk will be directly proportional to the distance to the router from the CCL.

It also may be tempting for users, once the NCP is connected to the aggregate routers, to switch away from the expensive SDLC, Frame-Relay, or X.25 communication lines or 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. Nevertheless, Visara can provide the upstream and downstream security application (SSL-1000) for encryption.

IBM: Enterprise Extender (EE)

zJournal Oct. 2001: CCL vs. EE

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 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.

Overview

Enterprise Extender is a feature that runs in conjunction with IBM's Communication Server for Linux (CSL). Like the CCL, this is a software only product, and therefore communicates through the OSA and similarly requires external hardware to provide any serial connections. Unlike CCL, EE is a relatively inexpensive software package from IBM and does not have the comparable licensing cost of NCP.

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 (Universal Datagram Protocol) encapsulation service that allows use of the IP infrastructure to carry the SNA frames to their destination. 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/High Performance Routing (HPR) is required, and the NCP functions will need to be disassembled and re-applied. This doesn't apply to the host applications, which remain ignorant 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.

Strengths:

- No NCP license
- Supports end-to-end IP
- Doesn't require migration of connections and interfaces, it replaces the entire network path.
- Been around for many years
- Is cleaner, more efficient and easier to support than CCL

Weakness:

- Migration to EE requires an EE endpoint on both sides of the IP network
- Requires APPN/HPR to replace SNI functions
- APPN/HPR must also reside on partner side
- Only z/OS supports EE; it cannot connect to CCL V1.2 or VTAM in z/VM, z/VSE or VSE/ESA
- Additional external router required

- Newer IBM CPUs do not support Token Ring OSAs, only Ethernet OSAs. To implement a Token Ring solution on newer mainframes would require enlisting an external router to provide the Ethernet to Token Ring conversion.
- Does not support EP/BSC

Enterprise Extender SNI Solution

CSL/EE does not support SNI, but it does support another feature that IBM markets as a replacement for SNI. This feature combination requires APPN/HPR traffic over TCP/IP. This feature along with the APPN Border Node function can provide an interface between two separate SNA networks and perform the same type of function that SNI provides. The main problem with this solution is that both networks (both companies) must have implemented APPN. Most companies will not implement APPN unless they absolutely have to, so unless both companies have already made the decision to run APPN, this solution is probably not workable.

Enterprise Extender X.25 QLLC Solution

CSL does not offer a direct solution for running X.25 from a mainframe Linux partition since the OSA is the only external communication route. The solution here would be to run SNA over LLC to an external router that would then have to provide an LLC to X.25 connection over a serial line. This solution is probably no worse and no better than running X.25 from the FEP-4600; it is converting ESCON/SNA to X.25 QLLC SNA.

Enterprise Extender Token Ring Solution

Support for CSL/EE Token Ring solutions is comparable to CCL Token Ring solutions. For the older CPUs that support Token Ring OSAs, the CSL Token Ring implementation is probably very straightforward. Newer IBM CPUs however do not support Token Ring OSAs, only Ethernet OSAs. To implement a Token Ring solution on these newer mainframes would require enlisting a router to provide the Ethernet to Token Ring conversion.

Comment:

The amount of effort involved with implementing APPN from scratch (man-hours) added to the costs of the product could offset the cost of our FEP-4600 solution. Our solution should be much easier to implement, saving the customer a lot of time and hassle. We do not offer any advantage in our Token Ring implementation if they have or can implement Token Ring OSAs. We do not offer much advantage for X.25 if the customer already has routers that can be reconfigured to support this. If the customer has to add routers for this function, then we might be an alternative. Our strength in face is if the customer has multiple needs, being able to provide a combination of features along with BSC 3270 would be a real selling feature.

Using Enterprise Extender to encapsulate SNA traffic across the IP network will be the natural migration from many mainframe customers. There remains a large portion of the market that requires very secure SNI traffic that may resist moving their data over an IP environment.

3745/46 Front-end Processors (Used/Existing)

The decision to add an additional 3745/46 or keep an existing one or choose to implement another solution is a big one. The most difficult hurdle to get over will be the old adage, "If it ain't broke don't fix it."

After announcing the EOL in 2002, IBM also announced that they would continue support for the 3745/46 family of products only until 2010. This means that many firms are looking for a replacement strategy for their existing products. While at the same time they are tasked with bringing up disaster recovery sites, and providing redundancy in their network as a result of the events of 9/11. IBM and their used hardware reseller present the market with the perception that the 3745/46 sets a rather high standard, and any replacement product will not do a good enough job. This perception is aimed at convincing firms to keeping the existing equipment or upgrading to the latest hardware set. There are a few large suppliers of 3745/46 hardware, ready and waiting to sell refurbished equipment and support it beyond the 2010 date. Some customers will hold onto their 3745/46 Front Ends while they hope or attempt to redo or replace the applications that they have been unable to migrate from.

3745/46 Overview

The 3745 is a Bus and Tag only platform. If you want to have an ESCON interface you must add a 3746. The 3746 is not a standalone product, so you will always see a 3746 as an add-on to a 3745. The 3745 runs licensed software called ACF/NCP (Advanced Communications Function/Network Control Program) usually referred to as just NCP. NCP pricing is based on the feature set of the 3745/46 (number of lines, etc.). The license is charged on a monthly basis. In addition to NCP, the customer will also have the need have ACF/SSP (System Services Program), another licensed software. This software is required to support NCP. If X.25 is also being used then there is a good possibility that they must also run yet another piece of licensed software, NPSI (NCP Packet Switching Interface). NPSI is required whenever the customer is running X.25 directly from their 3745. The 3746 expansion frame supports one flavor of X.25 QLLC (used for SNA traffic) that does not require NPSI. There is also a license associated with running EP (Non-SNA applications) lines on their 3745.

In addition to the 3745/46 Hardware and Software licensing (and maintenance), there is also the issue of NCP skill sets. The customer will either have to have someone on staff to perform the NCP responsibilities, or have to contract someone whenever changes have to be made. Because of the complexity of the 3745 with NCP, these skill sets are not easily learned or cheap.

Strengths:

- Main stay of corporate communications for over 30 years
- An untarnished reputation for reliability and service
- Provides stable and reliable communications for all of the traditional communications protocols (incl. SNI, Bisync, X.25, Token Ring, etc.)
- Very entrenched in the data center mainframe market

Weakness:

- Discontinued product with a limited support life
- Associated software licenses (i.e. NCP, SSP, EP, NPSI) can be excessive
- High maintenance fees

- Large footprint and environmental costs
- Device configuration for the 37XX is difficult and requires special skills
- Operationally, to implement changes to the network, the FEP must be shut down and the NCP software reloaded. This forces changes to be implemented during pre-scheduled off-hour times.

Comment:

The reputation of the 3745 is untarnished when it comes to reliability. In spite of the high cost of maintaining such a platform, most customers would prefer to keep such a platform, if it were not for the pending discontinued support for the product. The best strategy to displace the 3745/46 is on price. The 3745/46 platforms that we will be competing against will typically be Tier 1 or Tier 2 only (based on the number of lines). Higher Tier units will most likely require multiple FEP-4600s, a FEP-4600 with expansion chassis. Hardware platform cost configured with the number of lines that a FEP-4600 supports would probably be in the $20,000-40,000 price range, if bought from a broker. Tier 1 NCP licensing is about $415/mo, while Tier 2 is about $620/mo. To add ESCON capabilities, the customer will need to add a 3746 frame. This will add $450 or $900 per month depending on whether the 3746 has 1-4 processors, or 5 or more processors. To this amount you will need to add the tiered amount for NPSI if running X.25 (thought to be higher than NCP), EP (thought to be lower than NCP) if running BSC, and SSC.

Cisco Systems: Channel Attached Routers

Overview:

In the late 90s with the advent of the Internet and of IP as a universal protocol, an assumption had begun that there is no longer any need for FEPs as "everything will be IP". This has proved to be true at least in terms of where FEPs for the most part provide straight connectivity. Cisco was one of the first to recognize that most new mainframe-based application would be IP. Since this transformation began almost Cisco has captured most of the 90% percent of today's network access to existing and new applications that has moved to TCP/IP either natively or utilizing SNA/IP; most of which is via either Cisco's 7500/7200 channel-attached routers (now in End-of-Life), or their 6500 series Gigabit Ethernet-attached switch.

However, the remaining 10% FEPs perform a number of other vital communication related functions (SNI, Bisync, and others) that are closely linked to transaction applications including message and transaction switching, multiplexing, transaction security, QoS guarantors, and end-to-end transaction management and reporting. The need for these functions is especially important in mission critical transaction environments such as banking, government, point-of-sale, security, and health care applications. In these environments, FEP functionality is more necessary than ever before. Cisco does not have straight forward solution for these mission critical protocols.

An excerpt from a Cisco Systems "SNI replacement" white paper.

The IBM Communications Controller, commonly known as the front-end processor (FEP), has been used for many specific functions during the several decades that it served as a network attachment point for IBM mainframes. Of these, almost all either have been rendered obsolete or are now being done by more modern equipment. The only significant exception to this change is a function known as SNA Network Interconnection (SNI). SNI is

the connection, by gateways, of two or more independent SNA networks to allow communication between logical units (LUs) in those networks. In many cases FEPs that were purchased to support hundreds or thousands of networks devices are now being used to support a very small number of attachments to other SNA networks. This may create an extremely inefficient use of resources and a disproportionate cost to benefit. The largest impediment to replacing this technology is usually a result of it being embedded in at least two separate organizations, meaning that both must simultaneously agree to change.

The general advice from Cisco is to replace SNI at the application level. For each business function that is being performed through an SNI connection, there is almost certainly a readily available equivalent solution that uses the open IP protocol for connectivity.

Cisco acknowledges, however, that there are cases where an SNA-based alternative may be more expedient and, equally important, may require less inter-company agreement.

CIP 7500

The channel-attached 7500 series CIP (now in EOL) or 7200 Series with ECP4 routers or Gigabit Ethernet-attached switches (6500 Series with IBM OSA-E) are Cisco's primary product for migration from the 3745/46. Cisco routers do an excellent job of terminating SDLC or Frame Relay serial connections of any speed, and then passing that traffic across the channel to the mainframe. Like the 3745/46 they also support all standard LANs (using LLC2), including Token Ring and Ethernet.

The 7500/7200 series carries SDLC, Bisync and TCP/IP traffic over frame relay, or any of the other wide area network protocols. To connect to the mainframe it does not require software such as NCP, but does require the latest version of the VTAM operating system. It should be noted that it carries SDLC traffic only over frame relay connections and does not provide any direct connections for legacy SNA connections. When used to transport Bisync it must be connected to an IBM front end since its channel attachment is SNA only.

FEP-4600

Communications Controller Network Diagram