

Channel Interface Processor and Channel Port Adapter Migration White Paper

Document ID: 70636

Contents

Introduction

IP Datagram routing using CLAW or CMPC+

Feature Description

Suggested Alternatives

SNA – LLC bridging using CSNA

Feature Description

Suggested Alternatives

TN3270 Server processing

Feature Description

Suggested Alternatives

TCP/IP Offload

Feature Description



Suggested Alternatives

Summary

Related Information

Introduction

The Channel Interface Processor and Channel Port Adapters are widely used for network attachment to IBM (and plug compatible) mainframes and to provide services such as TN3270 conversion and TCP/IP Offload. Since Cisco has announced the End of Sale of these products, users of this equipment may want to begin planning alternative solutions, and this paper provides guidance in doing so.



To begin, it is important to note that there is no need to change immediately. There is adequate time to consider the options available to replace the functions of the CIP and CPA and to execute a migration strategy best suited to your situation. These are mature products that have been field tested in thousands of customer installations, encompassing tens of thousands of variations, and currently supporting millions of end users in production networks. Support for this equipment will remain available into the year 2011. We expect that for most customers, changes to their mainframe data center network should and will be driven by factors other than the eventual end of service of the Cisco mainframe channel products.

Over the last decade there have been huge changes in the design direction of mainframe networking. Plug compatible IBM mainframe vendors have left the market, allowing for a single unified approach to physical

network attachment of mainframes. The emphasis on traditional SNA subarea technology has been superseded by HPR SNA, especially to take advantage of HPR/IP and Branch Network Node capability. At the same time IBM has dramatically shifted their approach to networking on the mainframe, embracing an open systems model that maintains the same unparalleled level of availability required by the critical role of the mainframe in the enterprise. Ethernet Open Systems Adapters (OSA) with QDIO, and optimized for IP packet handling, provide a much more efficient path than ESCON channels to move data from the network to the mainframe. This foundation is then combined with Virtual IP Addresses (VIPA), dynamic routing protocols, and Quality of Service capabilities, to provide a complete foundation for high availability and high performance IP networking.


In most cases a new design that moves from the CIP and CPA to OSA includes an intelligent layer 3 switch such as the Catalyst 6000 with strong routing protocol and redistribution support and the capability of supporting a range of service modules.

IP Datagram routing using CLAW or CMPC+

This section provides information about the IP datagram routing feature of the CIP and CPA products.

Feature Description

Routing IP packets to mainframes was the first function to be implemented by Cisco CIPs, and Cisco's CLAW and CMPC+ channel protocols represent both the first and last channel protocols implemented on the CIP and CPA. They also represent the functionality most easily replaced, because the IP routing function is supported in all Cisco routers and layer 3 switches, and IP by its nature is independent of physical media considerations.



Suggested Alternatives

As the above diagrams show, the data center design may be simplified when using OSA interfaces directly attached to the aggregation layer in a data center. In either scenario, to provide maximum availability, a dynamic routing protocol should be run on the switch or router directly attached to the mainframe. The significant differences are that IP route aggregation is the primary function of the Aggregation layer switches, and they are designed to perform wire rate layer 3 switching, and serve as the control point for IP route redistribution.

This new design removes equipment that may incur costs for maintenance and operation, represents potential failure points, and introduces additional latency.

Assuming that the OSA interfaces are of the 100Mb Ethernet variety, and configured to operate in QDIO mode, they should provide similar, or slightly better throughput for IP datagrams than optimally configured (CMPC+ or CLAW PACKED) CIPs or CPAs, on a port by port basis. Obviously, for 1000Gb Ethernet, there is the potential for significant performance gains with the OSA design.


SNA – LLC bridging using CSNA

This section provides information about the Cisco SNA feature of the CIP and CPA products.

Feature Description

The CSNA feature provides bridging of LLC SNA traffic through a mainframe channel. Because of the variety of ways that SNA traffic is delivered to CSNA, the total solutions are generally more complex than those associated with IP routing. There may be any mix of local LAN attached SNA machines, DLSw+ delivering SNA traffic from remote locations, and SNA Switching Services (SNASw) routing SNA traffic using APPN. CIPs and CPAs running CSNA are also likely to be one of the few remaining places in a network where token ring technology is deployed, and a migration from CSNA should also include moving from token ring to Ethernet


A CIP or CPA installation for SNA may include any of the following elements.



Suggested Alternatives


Optimal conversion, SNASw used in branch routers

The simplest and most complete solution is to convert the existing layer 2 SNA traffic to use IP at layer 3 for transport, by connecting it to an SNASw router. If this is done adjacent to the layer 2 SNA machines, it limits the layer 2 SNA domain to small segments of the LAN, and removes any need to bridge this traffic across a WAN with DLSw, or between LANs.




Conversion to SNASw using DLSw+ in branch routers

An alternative solution, where it is not possible to install SNASw on the remote routers, is to use DLSw+ to bring the SNA traffic into the data center, and to then pass it off to SNASw for conversion into EE. While this still presents layer 2 SNA traffic in the data center, if the DLSw+ and SNASw features are both run in the same router, the layer 2 SNA will only be on a connection within those routers. Traffic arriving from the WAN and going to the mainframe will be IP.



LLC SNA bridged through Access layer to OSA in LCS mode

There are certain cases that require direct layer 2 connectivity between the SNA devices and the mainframe, and where IP based OSA-E is not useful. One such case may be where there are only local SNA machines and these require relatively high bandwidth connections to the mainframe. A second case is subarea host to host traffic that cannot be passed through SNASw and turned into EE traffic. Clearly this is the case especially for SNI or other traffic that is sent through an OSA to the Communication Controller for Linux (CCL) based NCP. You should consult the appropriate IBM documentation regarding configuring and managing OSA interfaces configured to handle LLC/SNA, or CDLC for CCL. For maximum performance and control you should try to place all of these SNA machines into one, or a small number of, layer 2 clusters within the access layer of the data center network. Token Ring attached devices present unique challenges, as not all data center infrastructure supports Token Ring attachment, and adding switches for Token Ring is very unlikely to be justifiable at this time. We suggest that Token Ring devices be attached directly to a branch router, and translational bridging be performed at that router. A form of redundant availability may be provided in the Ethernet environment by either of two methods. At the point that the SNA device attaches to the network, duplicate Ethernet MAC address may be used on a single LAN, with one of the address being suppressed until needed using HSRP. Alternatively, duplicate Ethernet MAC addresses may be used at the host end of the connection, by ensuring that these addresses exist on separate LANs, and that some form of spanning tree prevents them both from appearing on a common LAN.




TN3270 Server processing

This section provides information about the TN3270 Server Protocol feature of the CIP and CPA products.

Feature Description

The TN3270 Server is an industrial strength server, capable of reliably serving thousands of concurrent 3270 sessions. Its placement, as an integral part of the network infrastructure, provides design flexibility to achieve unparalleled availability.



Suggested Alternatives

We suggest that the only way to achieve similar scalability and availability is to place the TN3270 Server function directly on the mainframe. This provides a highly reliable environment, and with multiple interfaces and dynamic routing on the mainframe, continuous network availability. This also has the advantage of placing more of the complexity of SNA and its conversion to TN3270 into a single place, where the skill to administer it may be more readily available. There are two different mainframe based TN3270 Server program offerings available from IBM. The first is Communication Server (CS) for z/OS, included as part of the z/OS software. The other is part of the "Communications Server for Linux" offering.


TCP/IP Offload

This section provides information about the TCP/IP Offload feature of the CIP and CPA products.

Feature Description


TCP/IP Offload provides an alternative means of moving the payload data carried in IP datagrams across a mainframe channel. The objective is to handle some of the routine housekeeping duties of the TCP/IP protocol on the offload device, thereby lessening the amount of work required on the mainframe. While TCP/IP Offload was once widely used, efficiency improvements in mainframe handling of TCP/IP have largely eliminated the reasons for its use.

TCP/IP Offload Mode



- Mainframe TCP/IP cycle reduction

OSA-Express in QDIO Mode



- Dynamic routing for availability
- Wire speed performance
- Ideal place for Service Modules

Suggested Alternatives

For MVS systems using the IBM TCP/IP program, the decision whether to move from TCP/IP Offload has already been made, as support for offload ended at MVS Version 2.4.

Some customers are using the Unicenter TCPAccess Communications Server product from CA to take advantage of TCP/IP Offload. At an earlier point in time, this configuration represented the optimal performance model.. This product may also be part of a solution that provides TCP access to X.25 networks via X.25 over TCP (XOT). The simplest migration path is probably to change only those parts of the configuration that use the TCP/IP Offload function to use OSA–Express adapters instead. For those using other features of Unicenter TCPAccess Communications Server, this has the advantage of not disturbing those features. A more aggressive approach would be to consider changing the IP datagram access to use the IBM supplied stack, and if there are XOT features being used, investigating whether those could be enabled via the NPSI API interface to the CCL based NCP.

The TPF operating system has provided a full TCP stack, OSA–Express, and VIPA since 2000. It was originally enabled by PJ27333 in PUT 13 for TPF Version 4.1, and IBM reports dramatically improved performance and resource utilization using this model. While the TPF service model does not preclude customers from continuing to use TCP/IP Offload, we expect that the advantages of, and ease of moving to, the TCP/IP native stack support is compelling enough that TPF customers will want to change to this model prior to the end of TCP/IP Offload support.

Summary

Currently installed CIPs and CPAs will remain viable connectivity and TN3270 Server solutions for several more years. Beyond that, we expect that some quantity of CIPs and CPAs will continue to be available from refurbished stock. There are practical replacement solutions for each of the functions currently performed by the CIP and CPA. As an initial step, you should inventory the features and quantities of your current CIP and CPA usage. Then develop a plan to move, over the next several years, to a robust high speed intelligent layer three switch infrastructure to provide highly available and high–speed access to the mainframe.

Related Information

- **Cisco Channel Interface Processors**
- **Cisco Port Adapters**

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2014 – 2015 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jul 19, 2006

Document ID: 70636
