Third Party Peripheral Support

Patrick Hughes, Cray Research, A Silicon Graphics Company

ABSTRACT: *This paper details the third party vendor peripheral support offered by CRI. This paper contains a comprehensive list of peripherals supported by Cray. The process used to select, certify and qualify third party peripherals is described and the value added propositions provided by Cray are identified.*

Third party peripheral support

In most customers environments, the Cray system is a integral part of the customers computing enterprise.

As a part of the enterprise, the performance and reliability expectations of the Cray system and its peripherals are consistently the same as the other nodes in their enterprise.

The nodes are expected to be available, reliable and predictable. The goal of the computing enterprise is to provide a service that is as transparent as the water, gas and electrical services from the utilities.

To meet these expectations means Cray must be very cognizant of enabling technologies that can provide a competitive edge for customers.

Cray has for years maintained the strategy of supporting third party peripherals that meet or exceed market requirements and enable customers to have a competitive advantage in their industries or businesses.

While there are many vendor and technologies to choose from, large system customers continue to focus on a limited number of devices and interface technologies. Block / Mux and ESCON have been the channel interface defaults on the large Cray mainframes for the past 10 - 12 years.

When Cray introduced the YMP/EL series of computers and its successor, the J90, a whole new world of peripheral choices emerged. With access to the SCSI interface architecture, customers now have a wide variety of peripheral products and technologies at their disposal.

With the introduction of the Cray GigaRing architecture, J90se, T3E, and T90 customers now have access to peripherals that may not have been available in the past.

Although the list of third party peripherals supported on the Cray mainframes is very comprehensive and spans the price/performance spectrum, not every tape, disk, or network device that exists today will be found on the list (see Figure 1.).

Figure 1: THIRD PARTY PERIPHERALS SUPPORTED.

Vendor	Device	Interface	VME	Model- E	Giga-Ring
H/P*	C1533 DAT	SCSI-2	X		X
Exabyte*	8505	SCSI-2	X		X
CNT**	3200	SCSI-2	X		
Quantum	DLT 4000	SCSI-2	X		X
Quantum	DLT 7000	SCSI-2	1H97		1H97
STK	4280	SCSI-2	X		X
STK	4890	SCSI-2	X		X
STK	4781/ 4480	SCSI-2	X		X
STK	4791/ 4490	SCSI-2	X		X
IBM	3490E	SCSI-2	X		X
STK	9490	SCSI-2	X		X
STK	SD-3	SCSI-2	X		X
IBM	3590	SCSI-2	X		X
Ampex	DST-310	SCSI-2	1H97		1H97
IBM	3480	B/MUX		X	X
STK	4480	B/MUX		X	X
IBM	3490	B/MUX		X	X
STK	4490	B/MUX		X	X

Figure 1: THIRD PARTY PERIPHERALS SUPPORTED.

Vendor	Device	Interface	VME	Model- E	Giga-Ring
IBM	3490E	ESCON		X	X
STK	4490	ESCON		X	X
STK	9490	ESCON		X	X
STK	SD-3	ESCON		X	X
IBM	3590	ESCON		1H97	1H97
IBM	3494	n/a	X	X	X
STK	4400	n/a	X	X	X
STK	9310	n/a	X	X	X
STK	9360	n/a	X	X	X
STK	9710	n/a	X		X
EMass	ABBA	n/a		X	
E-Sys	ER-90	IPI-3		X	
Max. Strat	GEN-4	HIPPI	X	X	X
Max. Strat	GEN-5	HIPPI	X	X	X
NSC	N-130	FDDI		X	
NSC	PS-32	HIPPI	X	X	X
PSITECH	Frame-Buffer	n/a	X	X	X
Avaika	Switch	HIPPI	X	X	X
Netstar	Switch	HIPPI	X	X	X
FORE	VMA 200	ATM OC-3	X		
Interphase	VME 5215	ATM OC-3	X		
FORE	SBA 200	ATM OC-3		X	
Interphase	Sbus	ATM OC-3		X	
Interphase	V5211	FDDI	X		
CMC	CMC-130	Ether-Net	X		

Note: * 4mm DAT and 8mm EXABYTE tape device are not recommended as products to be configured for use on the J90 or any other Cray platforms when that use is for very intensive I/O functions such as DMF or other HSM applications.

Note: ** The CNT 3200 Gateway has been tested internally with STK 4480 tape drives only.

Cray adds value to the product

Before Cray announces its support for any product, several months of development effort has taken place to ensure product capability, reliability and performance meet the stringent requirements of Crays customers. These are the significant value propositions that are added by Cray:

1. Cray, with strategic input from our customers, maintains a strong influence on vendor future products and technologies.
2. Strategically, Cray stays on the leading edge with our partners developing products offering the highest bandwidth and capacities.
3. Crays early access to new and emerging technologies provides customers with products that positions them with a competitive advantage in their businesses.
4. Cray supported peripherals are considered an integral part of the system. Devices and their software drivers are optimized and tuned to take full advantage of mainframe, I/O controller and peripheral device capabilities.
5. Cray drives its I/O and devices very hard. With few exceptions, the extensive software and hardware testing programs for new devices have broke all new devices when they are first exposed to the Cray environment. This testing program results in a more stable and reliable product when it enters the customer environment.
6. All peripherals that are Cray supported are designed and packaged with serviceability and resiliency incorporated throughout.
7. Device micro-code testing and distribution is tightly controlled by Cray. This ensures that changes are tested prior to being loaded on customer devices. How many re-sellers and distributors even recognize changes to device micro-code, let alone test and distribute those changes?

How does Cray select vendors / products?

As technology evolves, product selection and transition becomes both easier and more difficult.

Easier as the density and bandwidth grow, the products and pricing become much more attractive to customers.

More difficult as vendor choices are also growing with both manufacturer and number of integrators grow. Additionally, as technology evolves, new recording techniques, faster rotational speeds, and smaller packaging become standard the options for new products become greater.

When vendors such as Cray make decisions as to what technologies and suppliers are being made, a number of factors must weigh in these important strategies and decisions.

An important ingredient in making any partnership arrangement with a peripheral manufacturer is ensuring the supplier and Cray share the same long term vision for the technology and products. In addition, the suppliers reputation for quality, on time delivery performance, and cost competitiveness are important factors in choosing the right partner.

From a product standpoint, reliability, capacity, bandwidth are some of the most important product selection issues. These, coupled with demands and market requirements are the driving factors behind product selection.

Cray maintains ongoing relationships with several vendors that have gained significant market share in their respective industries. As in selecting suppliers for Cray peripherals, the selection of third party product suppliers remains essentially the same.

While some mainframe suppliers may have adopted the strategy that device driver development and support is delegated to the device vendor, Cray maintains that these peripherals are an integral part of the overall system configuration.

Cray has taken the decision to invest heavily in the development of device drivers and the testing processes that go along with support for these peripherals.

Has every support decision been the correct decision? Absolutely not, there are a number of examples where Cray has acted upon the demand of a customer and has certified a product or technology that has come back to haunt both parties.

There are other examples where Cray has been slow to react to the unique demands and requirements of the marketplace and have caused un-necessary delays.

In some cases, customers have purchased unsupported devices, had their Cray service people install those devices and tried to execute. As unsupported and untested devices will normally do, they have failed. The results of these failures have been unplanned service interruptions, extraordinary expense for the customer and Cray to solve the problems

The certification and qualification process

Crays customers expect, and rightfully so, that the peripherals configured on their Cray systems function reliably and perform within the manufacturers specification.

While all hardware vendors truly have good intentions when it comes to delivering hardware, reality strikes home when devices enter the Cray environment.

Experience has shown that products advertised as ready for production use have consistently failed to meet even minimal performance and reliability goals when first exposed to Crays robust I/O drivers. Cray has consistently seen product testing cycles of 3 to 4 months and longer on some products.

As indicated in Figure 1, the bulk of the third party peripherals supported on Cray systems are magnetic tapes.

The next section of this document describes the testing process that is used for certification and qualification of all new tapes devices. This testing process is also executed when incremental upgrades to the IOS software are done and often when vendors ship pre-releases of device firmware to Cray for testing and certification.

This testing process is an extension of the comprehensive driver testing that is used on peripherals that carry the Cray name such as disk drives.

Phase 1 - Stand-alone

Basic Functionality Check

The focus is on verifying basic write/read capability, using the device in the mode for which the majority of use is expected-- file sizes ranging from one megabyte to 20 gigabytes, block sizes ranging from 64 kilobytes to 8 megabytes.

Extended Functionality Check.

The focus is on verifying complete data path functionality with UNICOS.

The following areas are tested:

I/O Types

- buffered i/o

system buffers used to block data during writes and reads

- unbuffered i/o

no system buffers used; only buffer is the user buffer

- transparent i/o

write(2) and read(2) system calls

- asynchronous i/o

writea(2) and reada(2) system calls

- tapelist i/o

use a list to describe a portion of the tape file

- Each list entry describes a data block, the presence of a tape mark, or the size of block

- flexible file i/o

use the FFIO library to provide record-oriented or byte stream-oriented data in an application transparent manner

- buffer i/o

- use the Fortran BUFFER OUT/ BUFFER IN interface

Multiple File/Volume Processing

- single-file, multi-volume processing

one file spans across two or more tape volumes

- multi-file, single-volume processing

multiple files written to a single tape volume

- multi-file, multi-volume processing

multiple files written across two or more tape volumes

File Processing

- appending to existing files

- concatenated I/O

dividing files across multiple devices

- ladder I/O

block sizes varied with every write

- multiple block size

block sizes consistent on a request, but varied with job

- position by file sequence

tape positioning using tpmnt(1) file sequence options

Positioning

- normal block positioning

use the ioctl(2) interface directly from C

- Fortran block positioning

use the Fortran settp(3F) interface

- absolute positioning

use the FFIO routine ffpos(3C)

- gettp processing

use the Fortran gettp(3F) call to obtain positioning data

Unix Utilities

- cp(1), cat(1), dd(1), cpio(1), tar(1)
- system level cpio and tar

perform actual file system archives

Fortran

- tsync(3F) processing

synchronize a tape, i.e. flush buffers

Tape Mark Processing

- write/read tape marks through list i/o
- write tape marks through Fortran endfile

Tape Label Verification

- ansi label, standard label, non-labeled, segy label type, bypass label

Blank Volume Processing

- write and read, load and unload a blank volume

End-of-Volume

- execute 20+ end-of-volume scenarios

Block sizes tested:

- 4096, 8192, 32768, 64000, 65536, 196608

Performance Verification

- verify write and read transfer rates over various block size
- verify timing on manual load/unload sequence

Phase 2 - Autoloader

Basic Functionality Check

- focus will be on verifying load/unload sequence
- execute jobs performing loads, basic writes and reads, and unloads for extended durations

Extended Functionality Check

- focus will be on verifying complete control and data path functionality with UNICOS. Areas covered will be the same as with manual drives

Performance Verification

- verify write and read transfer rates over various block sizes
- verify timing on autoloader load/unload sequence

Phase 3 - Autoloader with Mixed Media

Basic Functionality Check

- verify path to each media type

Extended Functionality Check

- execute jobs to mixed media concurrently.
- verify manual scenarios involving drive reselection, as well as execute functional tests from previous phases

By completing this process the devices are now ready for the enhanced functionality that is supported by the UNICOS tape daemon. This enhanced or advanced functionality takes the devices well beyond the default, character special functionality.

Character special support is defined as “the ability to open, to read from or write, and to close any device that is recognized by the I/O software.”

Advanced functionality includes all character special functionality plus features such as; dynamic resource control, dynamic configuration control, checkpoint / restart, loader domains, front end servicing, standard label support, multifile / multivolume support, imbedded tape marks, distributed operator control, end of volume processing, and absolute positioning.

In addition, support means that Cray and its support staff and support centers become very active in solving problems if they should occur. Cray becomes the single point of contact for the customer and there is a marked reduction in potential finger pointing between vendors when Cray is involved in problem definition and resolution.

Summary

As users tend to regard computing services much as they regard public utilities, the challenge of providing a reliable, available, and predictable level of service is becoming increasingly more difficult. As rapid as mainframe technology changes, peripheral device technologies such as magnetic tape and disk storage are on a more aggressive change cycle. There are many new technologies that claim to have all the answers, density, bandwidth, backward compatible, cost per megabyte and all those meaningful factors that work to influence product and operational strategies.

At the end of the day, decision makers are left with mountains of data and statistics that may enable them to make the correct purchase decision.

Cray brings value to our partnership with the customer by offering support for those products that are most widely accepted by the market. This support includes qualifying and certifying device functionality, going beyond basic character special functionality to a more comprehensive level.

Ongoing device support is carried through by testing as new releases of operating system and I/O software are generated. This support will most often be in place throughout the life-cycle of the device.

Cray works hand in hand with the customers ensuring that our products and those devices that we support contribute to your success in meeting your goals.

Tape technology is in transition

User requirements have outstripped tape technology. A huge speed and price gap exists between backup devices for mainframes and server environments. As client/server applications and network data proliferated, tape drive manufacturers were not positioned to address this technological urgency in a timely manner.

High- and low-end drive manufacturers were waiting on the sidelines to see if demand for mid-range storage would materialize.

Savvy users were demanding increased functionality, reliability and advanced data strategies, while vendors rushed to catch up with product solutions that only partially addressed drivers.

Consequently, these needs and requirements are being addressed by several storage vendors who are developing new tape products to fill this gap. For example, IBM has restarted its NTP2 or “coyote” project, StorageTek has recently funded its “eagle” project and Hitachi is predicted to have a new tape product within two years.

Meanwhile, other vendors such as Philips with NCTP, are tweaking their technologies, allowing backward compatibility and reuse of media.

Nonetheless, complete solutions, including storage automation for all enterprise tape operations, are some years away.

The problems and challenges will get worse before they get better. During the next two to five years, there will an explosion in tape technologies and products. Many of these products involve helical-scan recording technology. There is a lot of controversy surrounding this technology, but several commercial customers, including Cray customers have committed to and are embracing it. Meanwhile, traditional linear recording is being advanced threefold to twelvefold during the next three years. For the most part, the two market leaders and drivers will be IBM and StorageTek, although Exabyte with its Mammoth product has begun to show signs of becoming a new player in high-end tape. Quantum’s DLT is emerging as a leader in mid-range tape but does not yet play in traditional high-end markets.

Cray’s strategy for third party device support positions Cray and its customers well as magnetic tape technology evolves. Complementing this strategy are the strategic partnerships established with the mainline vendors over the years