Subject: Correct SPP Q&A

A living document. Now more correct than ever!

These are rough notes. MOST of this has been reviewed, but there
is some technical aspects that I am not sure I have exactly right.
Of course, there is more to find out, like scheduling.
Two caveats:

1) things change, usually for the better, and
2) I do not speak with authority, only with one-on-one conversations with worker-bees. No policy positions are to be contrued from this.

Timothy Barney 3/22/94

01) How far can teststation be from SPP?
A 50 foot cable come swith the system. It's a Thinnet cable
for ethernet physical layer protocol. There is a special protocol
upon that layer. So it is NOT compatible with regular ethernet.
This is a special cable in the regard of the protocol that it carries.
>>> You can NOT run ethernet protocol (other workstations) over this
cable. It is a dedicated cable between the teststation and the complex.

The purpose of the Test station is for service, and it should be
close to the actual system. Convex will only guarantee support for
the supplied cable between the teststation and the SPP.

02) How long is the FDDI cable that comes with the system?
10 meters. Most people were not aware that we wer shipping a cable
But I gather that we are FOR Preferred Partners as someone pointed
a stack of ~6 cables out to me that were intended as such.

03) What requirements are there for a network connection?
For the P3 systems ONLY, the teststations will be shipped with a
ethernet interface (thinnet and standard transceiver) so that the
teststation will accessible over local networks.
This is for Convex to access the teststation
over the Internet for support work on the P3 program.

The SPP will have a FDDI interface with aforementioned 10m cable.
It should be attached to an FDDI concentrator. It Can NOT connect
to a DUAL FDDI ring. The FDDI interface CAN connect to a single
attached ring, or a concentrator. The concentrator can then
attach to a dual ring. Customers must supply their OWN concentrator
if they want to attach to a dual ring.

The teststation and the SPP will EACH require a network address.
Again, Preferred Partners ALONE will be shipped and EXTRA ethernet
card for network access, besides the standard one for the DaRT
(diagnostic) bus. Regular shipments will not include this.

04) Other communications requirements?
A phone line for conversation will be REQUIRED.

A 9600 baud modem will ship with the DOMESTIC teststations.
International systems will be supplied a modem from the local
Convex office.
A second phone line for this is optional if Convex can access the
teststation over the Internet. Otherwise, it is required.

There will soon be a feature of CXTS such that when a certain level
of error is detected, the local support person is paged. For this,
the second phone line (for modem use) is a requirement.
The system administrator will define the level of error that pages
someone.

05) Can system administration be automated?

The system files on the teststation and SPP are standard system files of HP/UX. /etc/passwd, /etc/hosts and the rest. So any scripts you have that add batches of users will work.

06) Can I use my own HP for the teststation?
Not for P3 systems.
There are discussions as to offerring this option, but no decisions yet. There are technical considerations that may limit which workstations are appropiate for use.

07) How does the console function work?
At teststation boot time a console daemon is started. This will open an initial console window on the teststation. If you kill it with the X window manager (kill client) or exit the program driving it, the window will be restarted automatically after a specified time. This time is specified in /etc/inetd.conf. The -t option specifies the number of seconds. This may be changed to minutes. If there is no -t option, then the window is not restarted.

Alternately, you can rlogin into the teststation by means of opening a window locally or from a remote system (for example, your workstation in your cubicle). Once on the teststation, execute the program sn_cnsl with one of 2 options:

sn_cnsl -f 1	Force a console connection. This takes it away from the current console if one exists already. Only one console window may exist at a time (that has keyboard activated)
sn_cnsl -s 1	This opens a console connection in spy mode, where you may spy on the current console window. What happens if you do this and there is no current window? I don't know. Try it!

To exit fro this program type in CONTROL-E, then 'c', then '.'
^Ec.
To force a change in a spy window into the console window,
^Ecf

08) How do I boot the system (this changes/improves each week it seems)?
There are two ways to boot the P3 systems due to they're being alpha-release systems and still in development. The objectives are :
1. load server,kernel,tunables into memory
2. set up CPU environment (registers, etc.),start execution

Booting from SPP disk does both. Booting from teststation disk does 1. but not 2. So there is a way for OBP to do 2.

Booting from the teststation will NOT be supported. It has been a temporary way to boot. Now
that we can boot from SPP disk, it is on its way out. The capablity may well exist for initial P3 shipments but not for systems shipped in June I'm sure.

At reset, the primary loader is executed from MU EPROM. The Primary loader runs code on the CPUs in parallel to init cache and RAM. Primary loader loads the secondary loader, OBP (Open Boot Prom). OBP has 512KB of the 1MB of EPROM and uses 330KB of it. OBP, being loaded in memory, is entered on all CPUs in the node and one of the CPUs gets selected at random to run OBP and the rest go into an idle loop. Experiance shows this is usually/always the second CPU on the system. For sales types who hate the idea of selling a system that spends time being idle, call it a ready loop. The master CPU (randomly chosen) running OBP has the "autoboot?" configuration parameter turned off (for P3). So OBP prints it's FORTH prompt "ok" to console and sits,

waiting for input.

At some time (soon I think) this (autoboot?) parameter will be turned on and OBP will boot the system automatically from power on. Certainly this is how production will run.

When you execute the teststation command 'do-reset' to reset the system, the SPP goes through the above steps of a power on leaving you with the ok> prompt.

Here there are 2 options for P3 systems. Again, the capability to boot from teststation may not exist March 30.
After powering on the system (with the teststation utility pce_util or with main breakers) the system will boot up to OBP and stop with the "ok" prompt in the console window.

You can load the OS from the teststation disks or from SPP disks. From a teststation window (not the console window!) you cd to /spp/os and execute one of 2 scripts:

```
bootos          - load from teststation
boot.disk       - load from SPP disk
```

- bootos uses teststation utilities to take temporary control of the console, calls the script 'pushos' to push the OS (tunables, server,kernel) over to the SPP, then issues OBP command 'gohack' which sets the CPU environment and starts execution.

- boot.disk takes temporary control of console, then executes the OBP command 'boot', which is the same as OBP commands 'load' and 'go'.

In a window on the teststation, do:

```
/spp/os/bootos
```

It takes about 5 minutes to load the server and microkernel into memory from the teststation.

Alternately, from a teststation window you can do:

```
/spp/os/boot.disk
```

This takes less than 30 seconds.

It may be changed that to boot from SPP disk, you power on SPP and from console, where OBP 'ok' prompt will be, you just type 'boot'.

09) What is the LCD display used for?

On the 4 line LCD display, there are 4 lines. lines:
1. displays the node number and serial number
2. displays MU board use (error codes, etc)
3. OS info: 8 '-'s, with '-' being replaced with a '+' for each CPU enabled, and those '+' flipflopping with '*' each 1/2 second as a heartbeat.
   Then there are 8 '-' for each CPU that are not in use.
   This second set of '-' are replaced by CPU status for those CPU's active indicating what they're executing:
   'S' server code, 'M' microkernel code, 'I' idle,
   'U' user code.
4. OS info: Miscellaneous information from kernel

10) Can you explain the I/O subsystem -
the current implementation and the finished?

Yes.

11) Would you elaborate on that last answer?

Short of repeating the architecture manual, here are some fun facts.

The memory interface (MAUI, which is MU Board Adapter Interface, with A and U interchanged for a pronounceable name of a Hawaiin island) puts out a Cerebus to the Landfill board. This bus is multiplexed into two M Buses. These each drive two MBus-to-SBus chips. Each of these four chips drives a S bus. Attached to each S bus is one connector you see on top of the SPP node chassis. A FDDI interface card is attached to one of these (hardcoded as the third for now) and a SCSI interface card (hardcoded as first, Sbus #0 for P3).

The differential cable from the SCSI interface card to the disk chassis can go up to 25 meters. The cable shipping with the systems will not be 25 meters! The SCSI is Fast and Wide. At the disk chassis the differential is changed into single-ended SCSI which then goes to the disks. The disks have special small cards on them that when plugged into the 2 backplanes of the disk chassis identify to the disk what unit it is, so no jumpers need to be set.

The Sbus connectors are numbered:

```
Fan end of node                  Front of node
        U0SB0   U0SB1   U1SB0   U1SB1
          \                \___FDDI interface
           \_____ SCSI controller (P3 gets one for the 2 disks)
```

For Production, the other two Sbus connectors are for more SCSI controllers.

12) Can I move disks from my HP workstation to the SPP and vice versa?

No.

13) How do I disable and enable CPUs and memory slices?

First, a slice is a memory board, memory controller, CPU agent, and the two associated CPUs.

You may run on 2,3, or 4 slices.
You may run with 1-8 CPUs.

Each slice provides units of 4MB of continguous memory. You need 8MB of continguous memory for OBP and kernel to reside in, hence at least 2 slices.

On the teststation, there are utilities to configure the SPP.
'xconfig' is a X applications with point and click. Nice.
'ccmu' is a more flexible command line based program. This has a 'regen' command to scan the hardware present and auto-config the system based on the findings (# CPUs, etc).

regen is now run automatically when the node is powered on. After applying power, wait about 30 seconds to insure it's completed.

14) What do the 4 red LEDs on the MU board (middle, utility board) mean?

From LEFT to RIGHT:

- 68040 busy
- DaRT bus Active
- Undefined
- Heart beat of MU board

15) What is this scm thingy?

SCM is Sub Complex Manager. There is a GUI interface and an ascii interface. You allocate memory (in 16 MB chunks) and CPUs to sub-complexes.

Some info is available on the SPP by executing /usr/contrib/etc/ascii_scm without arguments - this lists the 26 commands.

16) Fixed priority - what is that?

Let me explain scheduling a bit.

A process.... To BE Written

17) Is there a limit to number of people able to rlogin to the SPP?

Apparently this is configurable; I need to contact someone about this.

18) Is it true that one teststation can service multiple SPPs ?

No.
A SPP will NOT boot and run without a teststation.

The teststation provides all environment monitoring.

If there are problems with cooling, or power, damage could result to hardware instead of being flagged earlier. Developing problems are logged to an event log that is on the teststation.

19) I've got my SPP. Now how do I log in?

There is a user sppuser with password "spp user" on the teststation. The root password for test station and SPP will be given to you by the install team.

20) How are CPUs numbered?

The hardware numbers them 0-7, by assigned slots.
So if only 4 CPUs are present, they should be numbered 0,2,4,6.
If all eight are present, then how does SCM number them?
Well, it depends on which ones are enabled? SCM numbers them according to the order it finds them. If 8 CPUs are present, and 0,2,4,6 are enabled and 1,3,5,7 are DISabled, then SCM will refer to CPUs 0,2,4,6 are 0,1,2,3 respectively.