# The Relative Costs of Providing CPU Standard Hours

-ooOoo-
J.J.Bunn, L.M.Robertson
-oOo-
11–March–1991

The CERN Computer Centre currently houses large machines offering different computing services to the community. In this paper we evaluate the relative costs to CN Division of running, maintaining and upgrading the three most important services, namely the IBM service, the VXCERN service and the Cray service. **This in no way takes into account the relative VALUE of each service to the user community**. It must be remembered that these three services have quite different flavours:

- IBM : Batch production, Interactive, General Infrastructure
- VAX : Interactive, General Infrastructure
- Cray : Specialised Physics Production

We thus calculate figures for the cost per interactive user, the cost per seat and the cost per hot seat for the services, which are in some way linked with their respective value.

We have then made a realistic estimate, along the same lines, of providing a batch service on the *shift* facility.

The purpose of this paper is **not** to show that *shift* is cheaper than central mainframe systems, but rather to highlight areas in which there may be cost benefits with a *shift*- style facility.

It is worth noting that the figures presented represent a *snapshot* in time, and would be quite different if, for example, any of the services were upgraded at no increase in cost.

## Existing Services

Table 1 shows the capacity of each main service, together with the delivered fraction of this capacity in 1990.

**Table 1: The Main Central Services**

| Service | Total Power in CERN Units | % Delivered in 1990[1] |
|---|---|---|
| IBM & Siemens | 42 | 90.5 |
| VAX 8800 & 8650 | 4.5 | 61.8 |
| Cray X-MP 48 | 32 | 74.1 |

[1]The "% Delivered" column shows the average delivered capacity as a percentage of full capacity for each service in 1990.

In order to obtain a cost in terms of operations, expenditure (on equipment and software etc.) and personnel for each of these services, we have adopted a simplistic approach.

## Amortization

We have noticed that, although there are large fluctuations from year to year, the *average* investment and operations expenditure by the Division on each service is largely independent of the number of years averaged over. For example, the average expenditure on the

IBM/Siemens service is around 8 MSF whether one takes all years since 1984 or just the last three years. To reduce the effects of amortization we have thus averaged the sum of the Divisional investment and operations expenditure over the last three years (1990, 1989 and 1988). The yearly costs and the resulting averages are shown in Table 2. Note that these figures are in agreement with those presented by D.O.Williams to the DPG at its zeroth meeting.

**Table 2: CN Investment and Operations Expenditure in kSF, 1988—1990**

| Service | 1988 | 1989 | 1990 | Average |
|---|---|---|---|---|
| IBM & Siemens | 7849 | 3570 | 12940 | 8120 |
| VAX 8800 & 8650[1] | 1787 | 1393 | 774 | 1318 |
| Cray X-MP 48 | 3691 | 4715 | 3730 | 4045 |

[1]Note that this figure does not include PRIAM

## Administrative Overheads

To adjust the operations plus expenditure costs for the overhead in administration etc. at CERN, we must add 19.5%. This figure was used by L.M.Robertson and A.Osborne in their 1987 paper "Cost of Computer Centre Services", which in turn was obtained from some CERN-wide calculations provided by J.Spalter. We believe that, if it was correct then, there is no reason it should be incorrect now!

## Direct Service Support

We have examined the list of staff engaged by the Division (but excluding visitors). From this list, we have removed, for the purposes of our calculation, the whole of AS, CE and CS groups, since those groups are not involved directly in support of the main services. We have also removed, for the same reason, two-thirds of DI group and those members of SW group involved in database support. We have then attributed the work of the remaining CN staff amongst various main activities, shown in Table 3.

**Table 3: CN Staff Distribution of Effort for Central Services - Full Time Equivalents**

| IBM | VAX | CRAY | PRIAM | WS | SHIFT | GENERAL+TAPES | Total |
|---|---|---|---|---|---|---|---|
| 12 | 7.5 | 4 | 1 | 8 | 4.5 | 24.5 | 61.5 |

In Table 3, the figure for GENERAL+TAPES includes a percentage of those people involved with administration of the Division as a whole. If we assume that those people, together with staff involved with tape support, are a necessary overhead for each of the supported services, then the service cost in terms of manpower in the Division may be expressed as:

$$servicecost = FTE * manyearcost * \{1 + 24.5/37.0\}$$

where the *man year cost* is currently 110 kSF to CERN, and FTE is the number of Full Time Equivalent staff attributed to the service.

Table 4 thus shows the personnel costs for the three main services:

**Table 4: Personnel Costs in kSF for the Central Services**

| Service | FTEs | Service Personnel Cost in kSF |
|---|---|---|
| IBM & Siemens | 12 | 2194 |
| VAX 8800 & 8650 | 7.5 | 1371 |
| Cray X-MP 48 | 4 | 731 |

## Administrative Overheads

Following the method in the Osborne/Robertson paper of 1987, we adjust the personnel costs for administration at CERN by adding 23%.

## Delivered Capacity

We now find a figure for the *delivered capacity in CERN Standard Unit Hours* for each of the services. A CERN Standard Unit Hour is equivalent to accounting the use of an IBM 168/E to one user for one hour.

$$Delivered = Availability * HoursInYear * CPUPower * 50/52$$

where 50/52 represents the fraction of the year that is accountable by COCOTIME. Table 5 shows the results.

**Table 5: Delivered Capacity in CERN Standard Unit Hours for the Central Services**

| Service | Delivered Standard Hours |
|---|---|
| IBM & Siemens | 320,000 |
| VAX 8800 & 8650 | 23,000 |
| Cray X-MP 48 | 200,000 |

With the above figures a cost for providing a *CERN Standard CPU Hour* for each major service is calculated. The figures are obtained by adding

- the Investment and Operations cost for the service, adjusted for CERN overheads, and
- the Personnel cost for the service, also adjusted for CERN overheads.

**Table 6: Cost of Supplying one Standard CPU Hour for each of the Central Services**

| Service | Investment and Operations Cost | Personnel Cost | Cost of one Standard Hour in SF |
|---|---|---|---|
| IBM & Siemens | 9703 | 2699 | 39 |
| VAX 8800 & 8650 | 1575 | 1686 | 142 |
| Cray X-MP 48 | 4834 | 945 | 29 |

## Service Values

As already pointed out, the relative costs presented in no way take into account the *value* of each service to the user community. One attempt to take the value of an interactive service into account is to derive a figure for the service cost *per interactive user* per year. This is defined as:

$$CostPerInteractiveUser = StandardHourCost * \left\{ \frac{DeliveredStandardHours}{UsersperWeek} \right\}$$

where the *Delivered Standard Hours* is as in Table 5 for the service. The *Users per Week* figure is an indication of the average number of *active* users that are liable to log in to the service at any time. (Those figures are obtained from the weekly statistics generated by CO Group.) The results are shown in Table 7.

**Table 7: Cost per Interactive User per Year in SF**

| Service | Standard Hour Cost in SF | Delivered Standard Hours | Users per Week | Cost per User in SF |
|---|---|---|---|---|
| IBM & Siemens | 39 | 320,000 | 3600 | 3,467 |
| VAX 8800 & 8650 | 142 | 23,000 | 1900 | 1,719 |
| Cray X-MP 48 | 29 | 200,000 | 100? | 58,000[1] |

[1]Note that the Cray is not an interactive service... see below

As already mentioned, the problem with this calculation is that we have three *types* of service, namely:

Batch production : IBM, Cray
Interactive : IBM, VAX
General Tape Manipulation : IBM

The numbers we have calculated in Table 6 assume each service offers just batch production to the user community; this is correct only in the case of Cray. The numbers in Table 7 assume interactive use (for batch production and tape manipulation the number of weekly users is irrelevant)—and so are correct only in the case of the VAX, where there is little real batch production achieved. (In particular, the Table 7 numbers do not take account of the roughly 60% non-interactive use of the IBM!)

So, we derive a figure that is the service cost *per seat* per year, defined as:

$$CostPerSeat = FractionToInteractive * \left\{ \frac{ServiceCostperYear}{UsersPerWeek} \right\}$$

where *Fraction To Interactive* is the fraction of the service that is accounted to interactive use. Table 8 shows the results:

**Table 8: Cost per Seat per Year**

| Service | Fraction to Interactive | Service Cost per Year in KSF | Users per Week | Cost per Seat in SF |
|---|---|---|---|---|
| IBM & Siemens | 0.28 | 12,480 | 3600 | 970 |
| VAX 8800 & 8650 | 0.74 | 2,840 | 1900 | 1110 |
| Cray X-MP 48 | 0.01? | 5,800 | 100? | 580? |

Our final table in this part of the paper shows the so-called *Cost per Hot Seat* for each service. This takes into account the actual number of interactive users normally logged in to each service at peak times, and is defined as:

$$CostPerHotSeat = FractionToInteractive * \left\{ \frac{ServiceCostperYear}{PeakUsers} \right\}$$

**Table 9: Cost per Hot Seat per Year**

| Service | Fraction to Interactive | Service Cost per Year in KSF | Peak Users | Cost per Hot Seat in SF |
|---|---|---|---|---|
| IBM & Siemens | 0.28 | 12,480 | 860 | 4,060 |
| VAX 8800 & 8650 | 0.74 | 2,840 | 200 | 10,500 |
| Cray X-MP 48 | 0.01? | 5,800 | 10? | 5,800? |

## The Cost of Providing Batch on Shift

Table 10 shows the prices of the *shift* equipment planned for the Pilot Project, assuming standard discounts. (Note that these are not the prices paid, as some of the equipment was donated, and some has been recycled.)

**Table 10: Standard Discounted Prices of the *shift* Equipment in kSF**

| UltraNet | CPU Servers | Disk Servers | Disks | Tape Units | Miscellaneous | Total |
|---|---|---|---|---|---|---|
| 624 | 640 | 249 | 104 | 92 | 100 | 1809 |

The maintenance of the equipment is estimated at 10% of the cost, and the costs themselves are amortized over two years. The total "operations and investment" cost of the equipment per year is thus

$$Operations/InvestmentCost = \left\{ \frac{1809}{2} \right\} + 0.1 * 1809 = 1085kSF$$

Following the "Direct Service Support" calculation above, we obtain a Service Cost for *shift* (assuming no manpower increase) of:

$$servicecost = 4.5 * 110000 * 1.662 = 823kSF$$

Both the operations/investment cost and the personnel cost must now be adjusted for administrative overheads. The expected availability of *shift* is around 80%, for a total CPU capacity of 100 CERN Units. Table 11 shows, for *shift*, the investment, operations and personnel costs, the expected delivered Standard Hours, and the resulting expected cost of one Standard Hour in SF.

**Table 11: Cost of Supplying one Standard CPU Hour on *shift* 20 GBytes 100 Units**

| Investment and Operations Cost | Personnel Cost | Delivered Standard Hours | Cost of one Standard Hour in SF |
|---|---|---|---|
| 1297 | 1012 | 674 000 | 3.4 |

Another *shift* configuration might include an increase in disk capacity to 150 GBytes. This would increase the Investment and Operations cost by 680 kSF (at $4k per GByte), and result in a Service Cost shown in Table 12:

**Table 12: Cost of Supplying one Standard CPU Hour on *shift* 150 GBytes 100 Units**

| Investment and Operations Cost | Personnel Cost | Delivered Standard Hours | Cost of one Standard Hour in SF |
|---|---|---|---|
| 1785 | 1012 | 674 000 | 4.1 |

With the increased disk capacity, the addition of CPU servers to bring the power of *shift* to 200 CERN Units would entail an outlay of 640 kSF, and result in the Service Cost shown in Table 13:

**Table 13: Cost of Supplying one Standard CPU Hour on *shift* 150 GBytes 200 Units**

| Investment and Operations Cost | Personnel Cost | Delivered Standard Hours | Cost of one Standard Hour in SF |
|---|---|---|---|
| 2243 | 1012 | 1 348 000 | 2.4 |

## Conclusions

We leave it to the DPG to draw its own conclusions from these figures. One major question that is as yet unanswered is: can one standard hour delivered on the IBM be meaningfully equivalenced to one standard hour delivered on *shift*?