AIM

Association for Information and Image Management

1100 Wayne Avenue, Suite 1100
Silver Spring, Maryland 20910
301/587-8202

Centimeter

Inches

MANUFACTURED TO AIM STANDARDS
BY APPLIED IMAGE, INC.

RHIC Off-line Computing

Study Group Interim Report

June 30, 1993

J. Featherly, B. Gibbard, J. Gould, M. LeVine, W. Love, S. Manly, T. Nepsee
M. O'Connor, S. Sorensen, D. Stampf, T. Throwe, D. Weygand

A report was prepared in Sept 1992, RHIC/DET Note 8, also known as ROCOCO, which estimated the various computing resources which will be required by the RHIC experimental program. A study has now been undertaken to review technical issues associated with supplying these resources. This study, organized by the HEP/NP Computing Group but including other appropriate participants, addresses questions of technologies, manpower, cost and schedule. The following document is an interim summary of this study both in terms of discussions which have occurred and initial conclusions reached.

• NEEDS SUMMARY

In the ROCOCO report one finds a summary of the estimated data processing requirements of the RHIC experiments. Between the preparation of the ROCOCO report and the current study, the estimated CPU needs increased dramatically on the basis of the reconstruction of Monte Carlo events using more realistic code. This increase from 40 GFLOPS to 160 GFLOPS is indicated below in Table I which is otherwise largely extracted from the ROCOCO report.

CPU	160	GFLOPS
Disk	4	TBytes
Raw Data Volume	300	TBytes/yr
Raw Recording Rate	50	MBytes/sec
Local Users	300	
Local Workstations	100	

Table I: RHIC Computing Load

• HARDWARE CONSIDERATIONS

Raw data is defined as that data which is recorded by an experiment while running for later analysis. The 300 TBytes indicated above is the raw data yield expected from an anticipate annual 2000 hours of running. The actual volume of data produced by an experiment is usually a factor of 2 or 3 larger than the raw data. Aside from raw data the other principle contributions are large Monte Carlo data sets and the first pass output of reconstruction results. In addition to originally writing the raw data, it is

MASTER

1 DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

necessary to read the raw data at least once and to write output results of comparable size in the course of Off-line processing. If the processing of data is to keep up with the taking of that data this leads to an Off-line mass storage I/O requirement that is a factor of 2 greater than the original recording rate. Thus in terms of the capabilities of the Off-line hardware the following Table is an appropriate target.

CPU	160	GFLOPS
Disk	4	TBytes
Total Data Volume	700	TBytes
Off-line Data Handling Rate	100	MBytes/sec

Table II: RHIC Off-line Computing Needs

1. Facilities for CPU intensive computing

The massive CPU intensive computing of a RHIC experiment will consist primarily of Monte Carlo simulations and the reconstruction of event data. There are three relatively distinct options for doing bulk CPU intensive processing. They are:

- (a) *"Farms" of independent workstations or workstation-like servers.*
- (b) *Commercially clustered sets of workstation-like machines, such as those in the most recent IBM offering, Scalable POWERparallel systems.*
- (c) *Symmetric multiprocessor systems where the CPU's actually share memory and peripherals such as in SGI's recent Power Challenge servers.*

In general, as one goes from systems of type (a) to type (c), the hardware itself becomes somewhat less cost effective, while the management of the systems becomes somewhat easier and thus the manpower costs associated with maintenance and operation are reduced. For all of these options there continues to be rapid technological advance, recently perhaps more so for the more tightly coupled processor systems. Software such as NQS, Condor, or Linda could be used with any of these options. CPU intensive technology is not thought to be an area in which there are especially long lead time considerations. The area is advancing so rapidly that it is difficult to engage in any kind of R&D which would give long term advantage relative to simply waiting for market developments. The best realistic price/performance ratios currently being seen are about \$400/MFLOPS. The rate of decrease in the cost is approximately a factor of 1.4 per year. This means that in 1997 a likely ratio is \$100/MFLOPS. This then implies that the cost of a system satisfying the 160 GFLOPS requirement in 1997 will cost approximately \$16M, well outside of the realistic funding profile.

Estimates of the short term computing needs of the experiments are not currently firm. Purchases in the style of the SGI and IBM systems already in place can probably cost effectively satisfy this demand in the short term while allowing for

hands on experience with up-to-date technologies and methods of deploying them. At a point about 2 years before turn on, a review of the state of technology at that time should be sufficient to produce a reasonably optimized system. The purchase of a significant fraction, 20-30% of the final system, would occur one year before turn on, to be followed by purchase of the remainder of the system the year of turn on.

2. Facilities for I/O intensive computing

It is not clear that this represents a different compute resource from that discussed above. However, it is the experience of other major experiments that in addition to the very CPU intensive operations of reconstruction and Monte Carlo, there are operations involving small amounts of computing on large numbers of large events. Such operations are often of the event filtering or event compression type. If the CPU intensive solution adopted above is a solution which lacks high speed bus access to the individual CPU's it could well be that that solution would not effectively deal with I/O intensive operations and that some other computing configuration with a higher I/O to CPU ratio will be required. It is imagined that this will again involve the use of standard commercial technology but perhaps of a somewhat different sort than that used above. For the time being we will assume that this means that a portion of the item above may be of a somewhat different form.

3. Data handling

The volumes of data and the data taking rates are comparable to SSC experiments but must be faced some years earlier. Many current HEP detector systems, dealing with very much smaller data sets, find that handling this data volume is the major computing bottleneck. It seems apparent that this is the area of greatest challenge for RHIC computing.

(a) Mass Storage:

In recent years many experiments have moved to 8 mm tape as the means of recording raw data because of its high storage density and the low cost of both the media and the tape drives. While 8 mm tape will probably have its place in the RHIC computing environment it seems unlikely that it will serve as the primary mass storage media. Its relatively slow recording rate and small volume size make it unattractive for large volume storage. On the other hand, 8 mm tape or, depending on the direction of product evolution, possibly Digital Linear tape or VHS tape may be used as a relatively inexpensive standard media for moving data samples to workstations and other small systems at remote institutions.

The following basic criteria apply to RHIC primary mass storage.

- *A single recording format should be used by all detectors for both On-line raw data recording and for Off-line large scale data storage.*
- *The recording system must be integrable into a large scale robotic system.*

- The recording rate for a single drive should be a significant fraction of the anticipated raw data recording rate of the large experiments.
- The capacity of individual media must be as large as practical.
- The cost must be as low as possible.

There are a wide variety of options which can be considered within present day recording technology. These include the following:

- 19 mm tape from Sony (D-1) or Ampex (D-2), recording rates to 32 MByte/sec, \$140 tapes hold 100 GBytes. Robots exist.
- VHS tape available from Metrum, recording rates to 2 MBytes/sec, \$20 tapes hold 20 GBytes. Robots exist.
- Digital Linear Tape, recording rates to 800 KBytes/sec, \$30 tapes hold 6 GBytes. Robots planned.
- Optical Worm with recording rates to 1.5 MBytes/sec, \$300 platters hold 9 GBytes. Robots exist.
- 8 mm tape from Exabyte, recording rates of 500 KByte/sec, \$5 tapes hold 5 GBytes. Robots exist.
- Creo optical tape, recording rates of 3 MByte/sec, \$10,000 tapes hold 1 TByte. No robots in sight.

In an effort to compare various current systems the following table was generated. Values reflect today's performance and costs while the standardized rates and sizes though somewhat arbitrary are approximately scaled to the RHIC needs specification. For general reference magnetic disk has been included while Optical Tape for which robots do not appear to be planned has been left out.

Media	Media Cost per GB	Cost of 300 TB of Media	I/O Cost per MB/sec	Cost of 100 MB/sec	Cost of 30 TB Robot	Latency (sec)
Disk	\$1,000	\$300,000K	\$ 300	\$ 30K	\$ 0K	.01
8 mm	1	300K	8,000	800K	3,000K	500
19 mm	2	600K	10,000	1,000K	400K	500
WORM	50	15,000K	14,000	1,400K	2,400K	20
VHS	1	300K	13,000	1,300K	1,000K	500
DLT	5	1,500K	6,000	600K	1,500K	500

Table III: Mass Storage Options

(b) File Serving:

In modern distributed computing systems where the operations are performed on a variety of network connect processors, code constants and data typically reside on file serving machines.

There are two important parameters which characterize a file serving system. The first is the volume of data it can hold. This in general scales directly with the cost of the system. The ROCOCO estimate of 4 TBytes of such random access served files represents a scale up of approximately an order of magnitude from present day HEP experiments while the scale up from present day experiments in terms of raw data is a full 2 orders of magnitude. While the purchase of disk at the 4 TByte level is likely to be feasible on the required time scale an order of magnitude more disk is not. The use of optical WORM technology which is cheaper by an order of magnitude per Byte turns out not to be practical because of the high volatility of the stored data and the relatively slow write speed.

The second parameter is the speed of access. In principle one would like to get the data as rapidly from the file server as from a directly connected disk. There are two major obstacles to this. First, there is the network and associated interfaces. Second, there is the potential for contention for any particular file or at least the disk that the file is on. The network and network interface problems are likely, as discussed below, to improve significantly over the next few years. The problem of real contention is not so obviously being addressed by advancing technology. There is remote access software such as AFS now (DFS in the future) which when used in place of NFS will cache served files locally on the reading machine when local disk space is available. This could help significantly for repeatedly accessed files of modest size. When there is a contention problem for one time accessed or very large files the only solution is to identify such files, make additional copies of them, and have a system which directs different users to different copies. An automated system capable of doing this should be sought.

(c) Advanced Hierarchical Data Serving:

In addressing the question of disk size, a natural approach to the desire to have a very large store of data on disk is to view the disk as a high speed cache for data in an even larger robotic store. This approach is being used with varying degrees of sophistication by existing HEP experiments. Putting together a system integrating high performance mass storage, large quantities of disk, and a general high speed network is clearly a very important RHIC project. Commercial systems based on UniTree and various combinations of computers, mass storage, disk arrays, and high speed interconnections are available on the market and perform automated hierarchical data management at the file level quite effectively. Other approaches are also being studied of which perhaps the most ambitious relevant work is in the PASS project intended to support SSC computing. This project utilizes an object database to permit direct access to exactly the event fragment of interest to a particular analysis. In its most interesting implementation such a system operates from a data store consisting of disk transparently serving as a cache for a very large, perhaps 19 mm tape based, mass store.

It should be pointed out that in many models of hierarchical storage, between disk, regarded as online storage, and tape, regarded as off-line storage, there is storage called near-line storage. A frequent medium for such storage is optical disk. It is characterized by a significantly shorter latency than tape but a cost which allows substantially greater capacity than magnetic disk. Limitations of such system include the fact that they have relatively slow I/O plus they are commonly "Write Once" systems while the data on them is rather volatile. These limitations lead to the conclusion that such near-line storage does not have a natural place in the RHIC data hierarchy.

4. Network

In recent years advances in CPU power have dramatically outstripped those in networking. The near term appearance of ATM and/or Fiber Channel as cost effective networking protocols can be expected to substantially correct this imbalance. There is also reason to hope that in the relatively near future network interfaces more akin to graphics interfaces will become commonplace reducing that bottleneck in network communication. These are again areas in which industry technology is proceeding at a pace which makes activity other than following the progress inappropriate. When decision time comes one will exploit the then available technology.

• SOFTWARE CONSIDERATIONS

As in the case of hardware there are various software options which may satisfy RHIC needs. Unlike hardware many of the decisions regarding these options do need to be taken very soon if one is to support the computing efforts of the experiments as they develop their software. Some of the options are discussed below.

1. Operating system options.

NextStep

MS-Windows NT

Taligent

Mach MK

VMS

"UNIX"

The choice of "UNIX" as the supported operating system is relatively straight forward. For now, support for additional, dramatically different, operating systems, such as VMS, seems out of the question on manpower grounds. However, within the UNIX world there clearly remain a large number of options (Berkeley, System V, OSF/1, etc). The expectation is that they are enough alike that multiple but not all versions of Unix can be supported. There are predictions that something like MS-Windows NT may sweep the computing world replacing Unix. Clearly if that happens adjustments will have to be made in terms of RHIC support.

2. General utilities

(a) Distributed computing options

DCE/DME (includes DFS, not from VMS)

SUN ONC, YP, NFS, DNS

AFS

Until recently it seemed likely that many vendors would adopt the DCE/DME standard associated with OSF. However recently a number of vendors have realigned forming an organization call Common Open Software Environment, COSE, which while stating goals similar to those of OSF is distinct from it and may define standards which are different. In point of fact such standards as the COSE members have so far agreed on are compatible with those of OSF so there is reason to believe that if there are two standard UNIX's they won't be too different. In any case it is probably not really necessary to select a distributed computing environment in a hard and fast way at this time. Clearly some products such as AFS, similar to DFS in DCE/DME, are important to doing the things RHIC computing needs to do. Such products should be introduced and used as needed without necessarily adopting an entire suite until such time as RHIC activities require such a decision.

(b) Graphics options

PHIGS

Open GL

GKS

X

Extensions to X (PEX, DPS, IEX, GLX)

The appropriate choice is not clear here. PEX which is the graphics extension to X seems like a good choice but is too low a level protocol to actually work in. PHIGS is intended as a primary user interface to this protocol and is currently available. However users seem to greatly prefer GL which is not yet generally available. While a GLX is now available for some display platforms it is also not generally available. The promise, in the near future, of a generally available open version of GL combined with wider distribution of GLX makes this seem like the long term best solution. Again this is an area in which it would probably be good to defer the final choice for a while.

(c) GUI options

Motif

Next Step

MS-Windows

Mac (Taligen)

Motif seems an appropriate choice at this time. It is widely used and there is substantial experience with it. There is the possibility that one of the others above may become dominant in the future at which point an adjustment might have to be made. Translation of an application from one Graphic User Interface to another is not out of the question.

3. Code development and management

(a) Code managements options

This is an area in which a very prompt choice should be made. There are currently two options under active use by RHIC experiments. STAR is using CVS/RCVS which is a Unix community/SLAC supported product while Phenix is using CMZ, a Patchy based commercial product distributed with the CERN library. The majority of knowledgeable people queried regarding this choice, recommend CVS or at least recommend avoiding CMZ. Phenix has agreed to try CVS and it is hoped that they will find it satisfactory so that support can be limited to it in the near future.

(b) OOP options

C++ seems to have become a defacto standard Object Oriented Programming language, at least in the Physics community. While experts frequently feel that some of the other OOP languages, Eiffel, etc maybe superior it is likely to be difficult to deflect enthusiasm from C++. At this point there does not seem to be pressure to make a choice but when there is it seems likely that C++ will win out.

(c) GUI builders

There are currently several GUI builders in use on site at BNL. One or more of them should be made available for immediate use. The most important property of any initial products is that they use standard User Interface Libraries so that the development work done with them can later be modified using other GUI tools.

(d) Case options

At this point no one within the RHIC or BNL community is working seriously with a CASE system. It seems clear that a CASE system of some kind should be acquired and made available for evaluation. Of the many available an Object Oriented system would be most consistent with recommended modern programming practices.

4. Major Applications

(a) Relational databases options

The selection of a relational database for use at RHIC is a major decision because of its anticipated wide use and expected large cost. The effort to define a set of criteria which can be used to evaluate various databases must be given high priority since the RHIC experiments already have database needs. There are a variety of uses for a database and it is not clear that a single database will satisfy them all. First, there is a good deal of administrative demand for databases (this may be best solved by PC databases independent of other demands). Second, there is a need for technical databases which describe the detector components, the various runs, and the various data sets all probably as a function of time. Third, there are more sophisticated database needs

such as the demand for reports in the form of specialized codes or routing tables which are used to actually configure the detector. Another sophisticated database application discussed earlier is the use of a database (probably an object database) to streamline physics analysis by allowing highly specialized access to the actual physics data. The mix in need between relatively small and inexpensive PC type databases and distributed client/server databases spanning an experiment is not yet clear. It is also not clear to what extent Object databases may satisfy some or all database demands better than Relational databases or how important various performance parameters are in a database. Prompt action with regard to databases is required.

(b) CERN library

The details of the RHIC CERN library support needs should be defined and BNL support for CERN library should be directly matched to those needs. This should be done as soon as possible.

(c) Other items

Editors, document preparation systems, presentation graphics systems, information handlers, etc, many of which are currently being used in the BNL HEP/NP community, should be more specifically specified in terms of RHIC needs and adjusted to satisfy them. These things need to be done sooner rather than later.

• SUPPORT

1. Platforms/Layered Support

A finite number of platforms can be supported with limited manpower. The expectation is that the finite number is approximately 3. Clearly if more manpower is available it is worth while to expand and support more systems, but the effective support of 3 platforms is much better than ineffective support of 6 platforms. At this time two platforms are being "fully" supported. They are SGI and IBM RS-6000. Some support is also being given to HP and there appears to be a significant need for support for SUN. The decision between HP and SUN as the other system currently being fully supported should be made as soon. There are at least three levels at which support is required. First, there is support for the basic hardware of a platform. At this time there are people and/or arrangements in place at BNL to assist in the installation and maintenance of the physical hardware of IBM, SUN and SGI. Second, there is need of support for the basic generic system software of platforms. Again for IBM, SUN and SGI support at some level of this type is available. Third, there is need for support for the particular environment and application that are used at RHIC for each platform. IBM and SGI are supported in a substantial way in this regard while minimal support also exists for HP, SUN and DEC/Ultrix. The support described is currently supplied by various agencies including CCD, the Laboratories HEP/NP Computing Group and contracts with outside vendors. It is expected that such a mix will continue.

2. Application Support

As a part of the selection and acquisition of an application it should be a goal that a person, generally the one responsible for the applications acquisition and installation, become sufficiently expert in its use to help beginners get started. It is expected that at sophisticated levels of use, users will become more expert than this assignee.

3. General Technical Support

An effort should be made to develop basic technical capability in areas which are likely to be of general interest so that appropriate initial support is readily available. Again it is anticipated that at sophisticated levels of usage, expertise beyond that of the user is not likely to be found locally. Developing a list of areas in which general technical support is important should be a first step in this process. The actual development of that support must clearly wait for the available manpower.

● STRAWMAN MODEL OF RHIC COMPUTING

The following is a perhaps somewhat arbitrary model of how computing might be done at RHIC on the basis of existing and projected technologies and costs. This then can serve as a test bed for the simulation and modeling of data handling and computing systems, which has been identified as a very important early activity in the RHIC off-line computing effort and as a reference to which other computing models can be compared. The real decisions on most of the individual components and technologies will not be made for at least 2 years and in some instances longer.

Assume that the technologies which will be used looks as follows:

1. CPU Intensive:

Commercially clustered 400 MFLOPS workstation-like servers at \$100 per MFLOPS. (projecting X4 cost reduction & no cluster cost)

2. Disk:

Raid, SCSI-2 connected via small, 100 GByte, servers capable of 16 MBytes/sec for \$400 per GByte. (projecting X2.5 cost reduction & no server cost)

3. Mass Storage:

19 mm at 16 MBytes/sec through small servers at \$150K per drive. (projecting no server or software cost)

4. Storage Robots:

2 TByte System, 1 Drive capable, at \$ 50K. 60 TByte System, 3 Drive capable, at \$400K. (projecting X2 capacity increase)

5. Network:

ATM directly connectable to all central systems, non-blocking at a 16 MBytes/sec interface rate, for \$3K per node. Expectation is that one node will handle 400 MFLOPS, 100 Gbytes, or 1 Tape drive. (projecting X8 cost reduction)

In this model one imagines the resources distributed in four distinct "clusters", One CPU intensive "production cluster" for each of the two large experiments, one CPU intensive "production cluster" shared by the smaller experiments and one "I/O intensive cluster" shared by all. Closely coupled to each of the CPU intensive "clusters" will be a small 19mm tape robot plus drive with sufficient disk to stage data from and to the robot for processing by that "production cluster". The I/O intensive "cluster" will contain the bulk of the disk and a large 19 mm tape robot equipped with 3 drives. The network will have sufficient bandwidth to allow the cross use of resources between these four "clusters". However, in general resources will be preconfigured to match allocations. The basic idea is that the production processing of data is done on the production clusters while general off-line access to the data will occur through the I/O intensive cluster. As an example the four "clusters" might at a particular time look as indicated in Table IV and shown in Figure 1 at the end of this document.

Phenix Production Cluster:

CPU	20	GFLOPS	\$2,000K
Disk	800	GBytes	320K
Tape System	1	Drive/Small Rbt	200K
Network	59	Nodes	177K

Star Production Cluster:

CPU	10	GFLOPS	1,000K
Disk	400	GBytes	160K
Tape System	1	Drive/Small Rbt	200K
Network	30	Nodes	90K

Small Exper. Production Cluster:

CPU	5	GFLOPS	500K
Disk	400	GBytes	160K
Tape System	1	Drive/Small Rbt	200K
Network	18	Nodes	54K

I/O Intensive Cluster:

CPU	5	GFLOPS	500K
Disk	2400	GBytes	960K
Tape System	3	Drive/Large Rbt	850K
Network	40	Nodes	120K

Total Hardware Cost:			\$7,491K
----------------------	--	--	----------

Table IV: Example "Cluster" Configurations

The above model fails to satisfy the Phenix CPU request while in all other areas it meets the original estimates. In Table V is shown the cost break down by type of resource. As can be seen, CPU costs already dominate this computing budget so no modest increase in the overall budget or redistribution of expenditures within it can come close to satisfying the stated Phenix CPU requirement.

CPU	40	GFLOPS	\$4,000K
Disk	4000	GBytes	1,600K
Tape Systems	4		1,450K
Network	147	Nodes	441K
			<hr/>
			\$7,491K

Table V: Computing Hardware Cost By Resource Type

Not included in this estimate are software and additional items such as printers, a spare 19 mm drive, workstations for the support staff and global network items which in total bring the overall capital cost of the RHIC off-line computing system to approximately \$8.5M in this model. Table VI contains a rough estimate of how capital expenditures might be distributed across years.

	FY '93	FY '94	FY '95	FY '96	FY '97
CPU			\$250K	\$1,000K	\$2,750K
Disk			100K	500K	900K
Tape Systems		\$150K	300K	500K	500K
Network			50K	200K	250K
Other		100K	200K	300K	450K
			<hr/>	<hr/>	<hr/>
			\$250K	\$900K	\$2,500K
					\$4,850K

Table VI: Capital Expenditures As a Function of Year

• SCHEDULE/MANPOWER

The overall manpower distribution shown in Table VII below is consistent with that produced by the formula in the ROCOCO document. It is intended to include all manpower required to setup and maintain the offline computing systems required by the RHIC experiments independent of the actual source of that manpower (RHIC hires, CCD, Physics Dept, etc). An important property of this estimate is that no provision has been made for a staff of operators to mount media, tear paper, etc. It has been assumed that the computing systems, including substantial robotic capability and

extensive user access, can be mostly self operated by what is expected to be a relatively sophisticated user community. It is also expected that data aide type support will also be supplied from within the experiments. To the extent that that is not possible there will be need for additional manpower.

	FY '93	FY '94	FY '95	FY '96	FY '97
System Support	2.0	3.0	5.5	9.5	12.0
Network Support	0.5	1.0	2.0	2.0	3.0
Code Devl. Support	1.0	2.0	3.0	4.5	5.0
Applic. Support	1.0	2.0	3.5	5.0	6.0
Technical Devl.	1.0	3.0	5.0	6.0	6.0
Hrdwr. Support.	0.5	1.0	2.0	3.0	4.0
Admin.& Manag.	0.0	1.0	2.0	2.0	3.0
Total	6	13	22	32	39

Table VII: Manpower for RHIC Computing

In Table VIII below, the manpower requirement is broken down in detail by technical activity. In Table IX, is shown an estimate of the effort for a few of what are expected to be the more significant projects of the RHIC computing effort. The manpower used in these projects is drawn from the various areas of technical activity listed in Table VIII. In these tables the symbols are used to indicate that the amount of effort is much less than an FTE, while symbols - - - - are used to indicate that the effort is about 1/2 of an FTE, simple integers are used to indicate greater levels of effort.

● CONCLUSIONS

Two aspects of satisfying the RHIC Off-line computing need have been discussed above. They are 1) the procurement and integration of hardware and software components into a system of appropriate capability, and 2) the formation of an organization technically capable of establishing, operation, and maintenance of this system. An additional aspect is the long term maintenance and evolution of this system throughout the life cycle of RHIC. This will include costs associated with the maintenance of hardware and software. The cost of such maintenance is typically in the range of 10-15% of the purchase price. In addition, given the fact that computing equipment becomes obsolete on time scale of 4 or 5 years, it will be necessary to replace the capital equipment at the rate of about 20-25% per year.

The above estimates and schedule description have been based on the assumption that RHIC will begin operation late in FY'97. Recent indications that a stretched out funding profile will result in a turn-on date in mid FY'99 will effect the schedule of off-line

	FY93	FY94	FY95	FY96	FY97
	:	:	:	:	:
General Systems Support	:	:	:	:	:
OS \& Environment Support	1111111	1111111	2222222	4444444	555
Platform Support	-----	1111111	2222222	3333333	333
Workstation Support	-----	1111111	1111111	2222222	333
X-terminal Support			-----	-----	111
	:	:	:	:	:
Network Support	:	:	:	:	:
Simulations					
Project Architecture		-----	-----	-----	---
Subnet Management	-----	-----	1111111	1111111	222
Data Server Integration			-----	-----	---
	:	:	:	:	:
Code Development Support	:	:	:	:	:
Code Management	-----	1111111	2222222	3333333	444
Object Oriented Language	-----	-----	-----	-----	---
CASE System		-----	-----	-----	111
Languages/Debuggers				-----	---
	:	:	:	:	:
Applications	:	:	:	:	:
CERN Library	-----	-----	1111111	1111111	111
GUI Builder			-----	-----	---
Database		1111111	1111111	1111111	111
Programming Graphics			-----	-----	---
Text Editor					...
Presentation Graphics					...
Documentation Preparation				-----	---
Information Handling/Access	-----	-----	-----	-----	---
Other				1111111	111
	:	:	:	:	:
Technical Development	:	:	:	:	:
Planning,Modeling,Integration	1111111	3333333	5555555	6666666	666
	:	:	:	:	:
Hardware Support	:	:	:	:	:
Workstations		-----	1111111	2222222	333
Servers	-----	-----	1111111	1111111	111
	:	:	:	:	:
Administration and Management	:	:	:	:	:
Admin. & Manag.		1111111	1111111	2222222	333

Table VIII: RHIC Computing Activities

	FY93	FY94	FY95	FY96	FY97
CPU Intensive Processing	:	:	:	:	:
	:	:	:	:	:
Simulations				...	
Prototype			2222222		
Initial System				3333333	
Full System					333
	:	:	:	:	:
I/O Intensive Processing	:	:	:	:	:
	:	:	:	:	:
Simulations				...	
Prototype			1111111		
Initial System				2222222	
Full System					222
	:	:	:	:	:
Data Store	:	:	:	:	:
	:	:	:	:	:
Simulations/Modeling	1111111	22222222	2222222	1111111	111
Prototypes		1111111	3333333		
Initial System				3333333	
Full System					333

Table IX: RHIC Computing Projects

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

computing activities and may in fact increase the possibility of satisfying the PHENIX request for substantially greater CPU capacity. While a detailed reconsideration of the schedule should now be undertaken some characteristics of the probable new schedule can be guessed. Since hardware, and to a lesser extent software, is becoming more cost effective each year, the purchase of the bulk of the computing equipment and hiring of that portion of the manpower required to support it, should be deferred until the two years immediately before the new turn-on date. However, since there is a great deal of leverage to be gained by early investments in planning, standards and software development, the ramp up in manpower of the first two years of the current schedule should be retained. The stretch out should thus result in a central plateau at about the midpoint of the staffing processes and a deferral in the capital equipment procurement schedule.

Figure 1: Example "Cluster" Configurations

**DATE
FILMED**

8/22/94

END