

## Statement of Need

### ***C.1 Procurement Objective***

The objective of this procurement is to provide the computational and associated resources necessary to support continued advances in environmental modeling capabilities and other high-performance computing system requirements that may arise within NOAA and at other partner agencies. This procurement will be known as the NOAA Research and Development High Performance Computer System (R&D HPCS) acquisition.

This procurement embodies the new approach that NOAA adopted to holistically manage its HPCS as a corporate asset. Prior to this procurement, NOAA organizations requiring HPC procured them along organizational lines. NOAA determined that its HPCS will be established and managed based on functional requirements. NOAA established two functional requirements to meet its HPC needs: (1) Operations and (2) Applied Research and Development. This Statement of Need describes the functional requirements of NOAA's Applied Research and Development.

A primary goal of the new approach is to achieve economies of scale in conducting its R&D HPCS acquisition while providing maximum flexibility in the resulting contract document. NOAA intends to achieve economies of scale primarily through consolidation of requirements into fewer acquisitions. In order to achieve maximum flexibility, NOAA's intention for its R&D HPCS contract is to be broad enough in scope and to have sufficient reserve capability (contract ceiling) to accommodate both existing and unanticipated requirements.

NOAA will use the Department of Commerce's acquisition process referred to as "Concept of Operations" or CONOPS, described in "Department of Commerce - Acquisition Process Case for Change" (available from the CONOPS home page located at <http://oamweb.osec.doc.gov/conops>). The intent of this process is to create an acquisition environment that will benefit the Contractor and the Government. In order to successfully implement this process within this acquisition, the Government seeks the cooperation of the vendor community in an effort to conduct business in an atmosphere of integrity, openness, and fairness. It is essential that the Government acquires the best HPCS available for the budgeted level of funding and to do so in an expedient and fair manner, therefore a R&D HPCS acquisition team ("the Project Team") was formed within the Department of Commerce. A Project Agreement was drafted between the team and management to spell out the objectives, milestones, approach, budget and resources available for the project. The Project Agreement and associated documents describing this project are available on the Internet at the Project Web Site, <http://rdhpcs.noaa.gov/>. Interested parties should monitor this web site for additional information over the entire lifecycle of this procurement project.

The contract supporting this acquisition consists of a four year base period, a four year option period, and a one year option to provide for contract transition (see Section C.9.1

for more information about the option periods). The Government expects that initial delivery, for Subsystems associated with workstreams 7-9, will occur in October 2005. The Government expects delivery, for any additional Subsystems or Components associated with workstreams 1-6, will occur in October 2006.

The Government requires a single Contractor to be responsible for the design, installation, maintenance, and support of the HPCS that will advance NOAA R&D's modeling programs. The HPCS shall meet the stated objectives and specifications set forth in this SON and shall include all hardware and software necessary to operate as a complete, functional, balanced, and highly reliable system. The HPCS comprises Subsystems that are designed to help accomplish NOAA's R&D missions. Some of the components that comprise a Subsystem include: Large-Scale Computing, Development, Post-Processing and Analysis Computing, Data Management, Hierarchical Storage Management, Interconnects, Software, and Visualization. A single Contractor will serve as the point-of-contact for the entire HPCS; even though the HPCS may involve components, from a number of different vendors, potentially located at multiple sites. Fundamentally, the Government must improve all aspects of NOAA's computing environment in order to fulfill its R&D mission. These aspects include:

- High-performance computing, including large scale computing, large scale data post-processing, analysis, and visualization capabilities. As described in this document and Section J of this RFP, the R&D workstreams are comprised of computationally intensive environmental modeling applications coupled to I/O-intensive codes and data storage.
- Hierarchical Storage Management System (HSMS). The HSMS shall provide archiving capacity to meet the expected rates of data production on the HPCS.
- Software for resource management, system administration, and application development. The Government requires operating systems and cluster software that can manage resources. Complete and functional FORTRAN90, C, and C++ application development environments shall also be provided.
- Reliability, availability, and support. Availability of at least 96% (24 hours/day, 7 days/week, calculated each month) has been the historical goal for NOAA's R&D HPCS. The HPCS must continue NOAA's historically high utilization of its computing resources. System reliability, availability, and Contractor support are considered fundamental aspects of the HPCS.
- High bandwidth connectivity, as required per workstream, to the NOAA Operational HPCS.

No less than 94% of the annual funding will be dedicated to the components of the HPCS specified in this RFP. Under direction from the Government, the remaining funds will be used to refine key areas of the HPCS or other aspects of the computing environment covered under the scope of the contract. Improvements will be made for performance, efficiency, or usability of the overall system. These areas may include, but will not be limited to, node, disk, memory, visualization, server, the network infrastructure, and additional support. The Government is not under obligation to direct this money to the Contractor. Key areas will be identified on an as-needed basis by performance assessments, including an annual system performance review by the

Government. The Government and the HPCS Contractor will work together to identify the necessary items that will best meet NOAA's computing needs. Actual purchases for this purpose will be determined solely by the Government.

The goal of this procurement is to acquire balanced, comprehensive computing capabilities to advance NOAA's research and development activities in environmental modeling. This includes not only high-performance computing but also additional hardware (storage devices and interconnects, communications interfaces, and other peripherals), software (primarily, but not limited to, the operating system, file system, storage management system, system utilities, performance and utilization software, diagnostic tools, and compilers), networking, full system maintenance, support services, and needed supporting infrastructure enhancements. This procurement will also support data management and archiving as well as analysis and visualization of model results.

The capabilities of the large-scale computers, hierarchical storage management system, analysis and visualization platforms, and network bandwidth shall be well-matched in a way that minimizes bottlenecks to the flow of information while maximizing performance. Achieving this proportionality in the acquired capabilities is an essential goal of this procurement. The Contractor is required to balance performance at all times.

The computing resources available to NOAA must meet its scientific needs throughout the life of the contract, so the Government requires a phased delivery of all components of the HPCS. The initial delivery of the HPCS must provide a substantial increase over current capabilities in computational throughput for NOAA R&D. At least one substantial upgrade to the sustained throughput must be provided during the base contract period, with archiving and other HPCS capabilities increasing commensurately. Individual components of the HPCS need not be upgraded simultaneously.

The basic tenets and provisions of this Statement of Need establish what the Government feels are the minimum acceptable capabilities of the HPCS based on NOAA's experience in performing its mission. However, innovation in proposed high-performance solutions is encouraged in addressing the needs of the Government. Newer technologies or an approach different from that presented here may provide opportunities to increase performance or enhance efficiency.

## ***C.2 Background and Purpose***

The mission of the National Oceanic and Atmospheric Administration (NOAA) is to understand and predict changes in the Earth's environment and conserve and manage coastal and marine resources to meet our Nations' economic, social, and environmental needs. NOAA's mission is embodied in its four strategic goals:

- Protect, restore, and manage the use of coastal and ocean resources through ecosystem-based management.
- Understand climate variability and change to enhance society's ability to plan and respond.
- Serve society's needs for weather and water information.

- Support the Nation's commerce with information for safe, efficient, and environmentally sound transportation.

In support of NOAA's mission and goals, NOAA conducts research and gathers data about the global oceans, atmosphere, space, and sun, and applies this knowledge to science and services that touch the lives of all Americans.

The ability to “predict changes in the Earth's environment” depends primarily on a diverse set of environmental models, usually requiring considerable computational resources. These models are developed through research and development (R&D) efforts within NOAA that occur primarily in three organizations: Office of Oceanic and Atmospheric Research's (OAR) Forecast Systems Laboratory (FSL), located in Boulder, Colorado; National Weather Service's (NWS) National Centers for Environmental Predictions (NCEP), located in Camp Springs, Maryland; and OAR's Geophysical Fluid Dynamics Laboratory (GFDL), located in Princeton, New Jersey. Computational modeling continues to emerge throughout NOAA and is germane to the daily operation of many of its laboratories. Organizations such as the Pacific Marine Environmental Laboratory (PMEL), Climate Diagnostics Center (CDC), Aeronomy Laboratory (AL) and Air Resources Laboratory (ARL) find using models for atmosphere, ocean, climate, air quality and ecosystem behavior to be instructive and informative.

### **C.3 Current NOAA R&D HPCS**

The current HPCS is comprised of three Subsystems with components that include an IBM Power 4 cluster; a Linux cluster based on Intel Xeon processors interconnected by Myrinet; and SGI Origin 3800, Origin 3900 and Altix systems. The entire HPCS includes over 7.5 PB of storage capacity, >15 million files, and over a Gb/s in WAN connectivity. See Appendix A for a more complete description of the current HPCS.

### **C.4 Benchmarks**

#### **C.4.1 Workstream (WS) Profile**

A workstream (WS) is a single instance of end-to-end processing (including pre- and post-processing). The set of instances of a given workstream within the Throughput Benchmark (see Section J) represents a set of jobs with similar characteristics within NOAA's current HPC infrastructure. NOAA's R&D HPCS Project Team defined 9 representative workstreams. See Section J for more complete descriptions of representative workstreams.

WS1	CM2-ESM	Coupled Earth System Model
WS2	CM2-HR	Coupled High Resolution Global Model
WS3	HIMF	Very High Resolution Ocean Model
WS4	EMTB	Environmental Modeling Test Bed
WS5	CMDC	Climate Model Development and Calibration
WS6	DAD	Data Assimilation Development
WS7	RUC	Rapid Update Cycle
WS8	WRF-CHEM	Weather Research and Forecast w/ atmospheric chemistry

WS9 WRF-SI Weather Research and Forecast w/ static initialization

#### **C.4.2 Benchmark Performance**

Within a workstream's funding profile (see Section C.4.3), the Government requires the maximum System Life Throughput (SLT) utilizing a balanced solution for each workstream in the contract period. The Government considers every workstream to be of equal importance and expects significant performance increases for all of them.

The benchmark is comprised of a set of application workstreams. Each workstream is a sequence of steps designed to represent the complete end-to-end execution of a single modeling application used in NOAA R&D. Each workstream may contain multiple applications with a mix of compute-bound and I/O-bound codes. The Contractor shall minimize the execution time of each workstream and maximize the overall throughput when multiple workstreams are run concurrently, as described in Section J.

SLT is the metric of system performance for each workstream. The workstream SLT is a measure of system performance and availability integrated over time. See Section C.6.1.2 for more on workstream SLT.

The workstreams were baselined on different computer systems. Namely, WS1 - WS3 were baselined on NOAA's Origin 3000 HPC in Princeton, NJ; WS4 - WS6 were baselined on NOAA's IBM Power 4 Regatta H HPC in Washington, D.C.; and WS7 - WS9 were baselined on NOAA's Xeon-based HPC cluster in Boulder, CO. Further details of these workstream baselines can be found in Section J.

The Government can most effectively use performance increases that are implemented at regular, but not too numerous intervals during the contract. Further, the Government will have difficulty using disproportionately high performance delivered either very early or late in the contract period. Accordingly, a given workstream should receive no more than two (2) performance upgrades during either the base or option periods.

The Government will evaluate performance increases in terms of the performance baseline described in Section J. The initial deliveries for WS7 - WS9 are required in Q1FY2006, the initial deliveries for WS1 - WS6 are required in Q1FY2007. The Government requires maximum System Life Throughput obtained by at least one significant mid-life upgrade during the base and option periods. It is desirable that the mid-contract upgrades for the base period for WS7 – WS9 occur between Q4FY2007 and Q1FY2008 while WS1 – WS6 upgrades should occur between Q2 and Q3FY2008.

Benchmarking for system upgrades shall be handled as follows:

- 1) Benchmark revision is required at the *option period upgrade* (see Section C.9.1 for information on the option period upgrade).
- 2) Benchmarks may change by mutual agreement between Government and Contractor at *mid-life upgrades*.
  - a. If a new benchmark is not defined, further changes to the benchmark software for performance will be allowed only at the discretion of the government.
  - b. Further, any software changes that are allowed will be used to create a new baseline with respect to the existing system.
  - c. Thus, a performance increment is always defined with respect to a constant set of source code across the change.
- 3) Revised workstream benchmark codes and performance baseline must be delivered to the Contractor no later than 6 months prior to delivery of an upgrade.
- 4) Performance data shall be reported as with the RFP structure:
  - a. Throughput and Scaling values with class A changes only.
  - b. Any additional Throughput and Scaling values with class B, C and D changes. All such changes will be evaluated for risk as with the RFP response.
- 5) The configuration management (CM) process will be utilized as the forum to discuss scientific advancement as it relates to model performance as well as software architecture changes to accommodate the new science. In the context of the site-specific CM meetings, the framework will keep both Government and Contractor personnel up to date with desired future trends and aid the benchmark revision process. Among other items, the CM forum will be used by the Contractor to suggest benchmark code changes to NOAA technical staff for technical evaluation.

### **C.4.3 Funding Profile**

Table I shows the projected funding levels associated with each workstream. The funding begins for workstreams 7-9 in FY2006. The funding for the remaining workstreams will begin in FY2007. The Government requires a reserve option where no less than 94% of the annual funding will be dedicated to the components of the NOAA R&D HPCS specified in the RFP. Assuming no change from FY2005 President's budget, it is anticipated that funding levels will remain constant for the duration of the base and option period (see Section C.9 for more information about contract options). Funding shall not cross workstreams, as shown in the table below. Based on information provided by the Contractor, the Government will evaluate proposals in order to verify that workstreams are appropriate to their funding profile.

The available funding level for the one year extension options is anticipated to be approximately half of the normal yearly expenditures. More information about the extension options can be found in Section C.9.

**Table I - Funding Profile in \$Millions (shown before 6% reserve is removed)**

		WS1	WS2	WS3	WS4	WS5	WS6	WS7	WS8	WS9	Total
Base	FY2006							\$1.3	\$1.0	\$0.7	\$3.00
	FY2007	\$4.0	\$5.7	\$3.8	\$1.48	\$1.48	\$1.48	\$1.3	\$1.0	\$0.7	\$20.94
	FY2008	\$4.0	\$5.7	\$3.8	\$1.48	\$1.48	\$1.48	\$1.3	\$1.0	\$0.7	\$20.94
	FY2009	\$4.0	\$5.7	\$3.8	\$1.48	\$1.48	\$1.48	\$1.3	\$1.0	\$0.7	\$20.94
Option	FY2010	\$4.0	\$5.7	\$3.8	\$1.48	\$1.48	\$1.48	\$1.3	\$1.0	\$0.7	\$20.94
	FY2011	\$4.0	\$5.7	\$3.8	\$1.48	\$1.48	\$1.48	\$1.3	\$1.0	\$0.7	\$20.94
	FY2012	\$4.0	\$5.7	\$3.8	\$1.48	\$1.48	\$1.48	\$1.3	\$1.0	\$0.7	\$20.94
	FY2013	\$4.0	\$5.7	\$3.8	\$1.48	\$1.48	\$1.48	\$1.3	\$1.0	\$0.7	\$20.94
<b>Total</b>		<b>\$28.00</b>	<b>\$39.90</b>	<b>\$26.60</b>	<b>\$10.36</b>	<b>\$10.36</b>	<b>\$10.36</b>	<b>\$10.40</b>	<b>\$8.00</b>	<b>\$5.60</b>	<b>\$149.58</b>

## **C.5 HPC Subsystem Components**

### **C.5.1 Computing Requirements**

#### **C.5.1.1 Large-Scale Computing (LSC) Component**

The Contractor shall provide a Large-Scale Computing (LSC) component at a substantial increase in sustained throughput over NOAA's current supercomputers described in Appendix A. Sustained throughput shall be measured by a throughput benchmark (see Section J) comprised of workstreams that are surrogates for NOAA's expected future workload. The metric of performance for the LSC is based on sustained throughput. The scalability of the computational platform(s) shall be measured by a benchmark designed to reveal the performance and scaling characteristics of individual codes as they are executed on different processor counts. The scalability measure will assist the Government in evaluating performance projections.

It is required that when any set of resources (such as a cluster node) in the LSC fails, batch jobs using those resources must be capable of being rerun without user intervention. In this situation, only interactive sessions hosted on the failed resources will be lost, and the Subsystem must allow users to continue to be able to login interactively. It is desirable that failover be to processors that are binary-compatible with and running the same operating system level as the failed processors. The capability of the LSC to operate and be repaired in degraded mode is required. It is desirable that the LSC have no single point of failure. The Government requires an availability level (defined in C.6.1.2) of at least 96% on every major component in the R&D HPCS.

At least one substantial upgrade to the sustained throughput of the LSC, as measured by a throughput benchmark, shall be provided during the base contract period. This has traditionally been considered the mid-life upgrade.

### **C.5.1.2 Development Component**

For each LSC Subsystem a binary compatible interactive development platform shall be provided. The Development environment shall have the same software environment as the LSC Subsystem. The interactive resource should represent 5% to 15% of the LSC component. It is desirable that the amount of interactive resource available is adjustable within the Resource Management Software (see Section C.5.3.1) without a system reboot.

The Government requires resources for interactive work. The Government desires that interactive resources have a minimal impact on the batch production resources and vice-versa. It has been NOAA's experience that the nature of interactive R&D work creates resource contention with batch production jobs, and batch production jobs slow interactive response time. The Government desires the ability to reassign batch production resources for interactive work without a reboot or a restart to the batch system.

The Government desires early access to the Development Component for training, software porting and tuning, and OS and application configuration testing. Early access (with some Contractor provided engineering assistance) up to 6 to 9 months in advance of delivery is desirable.

### **C.5.1.3 Post-Processing and Analysis Component**

Some of the HPC Subsystems described in Appendix A execute compute-bound and I/O-bound codes separately. For example, the climate computing resources, located in Princeton, use a Large Scale Cluster for the former and an Analysis Cluster, configured for very large sustained I/O, for the latter. The Contractor shall provide resources that can efficiently execute both types of codes, with the goal of minimizing the execution time of each workstream and maximizing the overall throughput when the workstreams are run concurrently.

## **C.5.2 Data Management Requirements**

The Government's scientific data are critical assets. As noted in section C.6, the Government requires the highest level of data integrity and at least 99% availability for access to its scientific data. The Government requires accessibility to its scientific data in the absence of the Large-Scale Computing resources.

### **C.5.2.1 Home File System (HFS)**

A Home File System (HFS) is required for each HPC Subsystem provided. This file system will have small quotas and be used to store source-code, small data sets, and environment initialization files. The Government desires initial HFS storage of several GBs per user with the ability to grow with each upgrade. The HFS shall be globally visible to the specific Subsystem with which it is associated. If the HFS is not available, no workload can be accomplished by the associated Subsystems. If no workload can be accomplished the Subsystem will be in downtime, as described in Section C.6.1. The HFS shall be backed up utilizing an automated method (see Section C.5.2.7).

### **C.5.2.2 Fast-Scratch File System**

A Fast-Scratch file system is high-bandwidth local storage that is visible by all processors within a given batch job. The Fast-Scratch file system will be purged frequently. The Fast-Scratch file system will not be backed up. Initially, per Subsystem, the Fast-Scratch file system shall be able to support 100 million files and individual files that are up to 500 GB in size. The Government expects that these requirements will grow over the life of the contract.

### **C.5.2.3 Long-term Scratch File System**

Long-term disk storage is a staging area for users to temporarily place files that will continuously be manipulated. Large data sets that are not packed into a larger file are not conducive to a HSMS file system. This file system does not need to be backed up and can be purged of aged data. Initially, per Subsystem, the Long-term Scratch file system shall be able to support 100 million files and individual files that are up to 500 GB in size. The Government expects that these requirements will grow over the life of the contract.

It is anticipated that the scientific data for WS4 – WS9 will be primarily accessible through the Long-term Scratch File System. The Government requires accessibility to its scientific data in the absence of the Large-Scale Computing resources.

### **C.5.2.4 Hierarchical Storage Management System**

The Hierarchical Storage Management System (HSMS) component will be purchased. The Government desires a flexible HSMS hardware configuration that can be changed or enhanced to optimize HSMS performance. Upgrades to the HSMS shall be provided commensurate with upgrades to the computational resources, with overall system balance the goal.

The Government requires accessibility to its scientific data in the absence of the Large-Scale Computing resources. It is anticipated that the scientific data for WS1-WS3 will be primarily accessed through the HSMS.

The Government requires a two-tiered storage scheme for its data archive, comprised of frequently accessed nearline storage (robotically mounted at high speed) and infrequently accessed offline storage (also robotically mounted), that will effectively satisfy the requests for scientific data that make up NOAA's R&D workload (as discussed in section C.4). The HSMS software shall provide automatic migration between data archive tiers based on a combination of access time and file size. The Government requires a daily report of usage by groups and users of all storage tiers. The HSMS shall meet the performance requirements of the HSMS Archive benchmark given in Section J.1.4.4.

NOAA's existing data archive for workstreams 1-3 resides on a Government owned, SGI/StorageTek HSMS in Princeton (see Appendix A). This legacy HSMS is available to the Contractor as unrestricted GFE beginning FY2007 (see

Appendix C). The Contractor shall make arrangements to keep this data available, potentially at a reduced level of performance. The Government desires a transparent migration of this legacy archive to the new HSMS, maintaining the present filesystem image.

NOAA's existing data archive for workstreams 7-9 resides on a Government owned, Sun/ADIC HSMS in Boulder (see Appendix A). This legacy HSMS will be under maintenance until 09/09/2006. At that point, the HSMS is available to the Contractor as unrestricted GFE (see Appendix C). The Contractor shall make arrangements to keep this data available, potentially at a reduced level of performance. The Government desires a transparent migration of this legacy archive to the new HSMS, maintaining the present filesystem image.

Multiple archives may be provided. If disk is required for the caching or staging of files, it shall be fault-tolerant and shall be provided in addition to the other file systems.

Each data archive shall be presented to the associated Subsystem as a single filesystem. Initially, per Subsystem, the HSMS shall be able to support 100 million files and individual files that are up to 500 GB in size.

The Government desires HSMS software comparable to San Diego Supercomputer Centers' (SDSC') Storage Resource Broker (SRB), which transparently combines logically related small files into one larger "container" file in the archive filesystem.

The Contractor shall provide all storage media used in the HSMS. The Government desires that files that haven't been accessed in one year will be migrated from the nearline storage to the offline storage. For nearline and offline data, the Contractor shall provide a data recovery service to recover data in the event of media failure.

If files are staged on disk, the Government desires enough disk space to keep files disk-resident at least 24 hours.

For the nearline tier, the Government desires a high file positioning rate,  $P$ , with matching robotic mount performance. For tape media,  $P$  (in files/hour) is given by the equation below (note that there is 3600 seconds/hour). The "average search time" used must apply to the method the HSMS software uses to position to files.

$$P = \frac{(n * 3600)}{(load + search + rewind + unload)}$$

- $P$  = file Positioning rate
- $n$  = Number of tape drives
- $load$  = time to Load a tape in sec.
- $search$  = average Search time in sec.
- $Rewind$  = average Rewind time in sec.
- $unload$  = time to Unload a tape in sec.

For the nearline tier, the Government desires the use of several types of archive media, optimized for the file size distribution, balancing the file positioning time and file read time.

The Government desires that files recalled from the offline tier automatically revert to nearline tier residency. User-specified migration between tiers is also desirable. A method for determining the physical location of users' files within the storage hierarchy is desirable.

The Government desires that the HSMS software provide priority scheduling of data recall requests, assigning high priority to certain users or to batch jobs over interactive sessions.

### C.5.2.5 Data Generation Profile

#### C.5.2.5.1 Data Generation for WS1 – WS3

NOAA's data generation for its computational baseline for WS1 – WS3 is shown below:

	WS1	WS2	WS3
Baseline Data Generation	1.4 TB/day	2.6 TB/day	1 TB/day

Data generation for WS1 – WS3 is related to the computational increases related to the workstream. Such changes shall be negotiated. Data generated ( $d$ ) by the workstream ( $i$ ) can be calculated using the following equation:

$$d_{i,j} = d_{i,(j-1)} \left( \frac{B_{i,(j-1)}}{B_{i,j}} \right)^{0.7}$$

$d$  = Data generated by the workstream  
 $B$  = workstream Benchmark time  
 $i$  = WS number  
 $j$  = system configuration period

Due to changes in the computational models, this equation may need to be changed throughout the life of the contract.

The following example demonstrates the application of the formula above. Assume that during the previous system configuration period ( $j-1$ ) the workstream was generating 1 TB of data per day and the workstream Benchmark time ran in 10,800 seconds. Also assume that during the current system configuration period ( $j$ ) the workstream Benchmark time now runs in 7,200 seconds. The data generation ( $d$ ) for the current system configuration period ( $j$ ) is given by:

$$d_{sample} = 1 \frac{TB}{day} \left( \frac{10800 \frac{s}{benchmark}}{7200 \frac{s}{benchmark}} \right)^{0.7} = 1.34 \frac{TB}{day}$$

### C.5.2.5.2 Data Generation for WS4 – WS9

Table II shows the projected data generation associated with WS4 - WS9 throughout the contract period.

Table II -- Data Generation for WS4 – WS9 (in TB/day)

	WS4	WS5	WS6	WS7	WS8	WS9	Totals
FY2006				2.4	2.2	1.6	6.2
FY2007	3	3	3	2.4	2.2	1.6	15.2
FY2008	5	3	5	4.8	4.4	3.2	25.4
FY2009	7	3	7	4.8	4.4	3.2	29.4
FY2010	10	6	10	9.5	8	6	49.5
FY2011	13	6	13	9.5	8	6	55.5
FY2012	20	6	20	12	10	8	76
FY2013	30	12	30	12	10	8	102
FY2014	30	12	30	12	10	8	102
<b>Totals</b>	<b>118</b>	<b>51</b>	<b>118</b>	<b>69.4</b>	<b>59.2</b>	<b>45.6</b>	

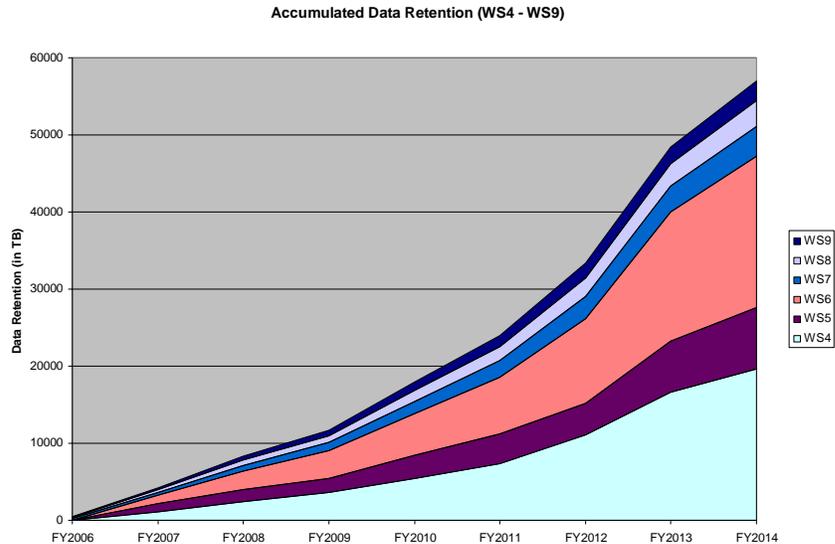
### C.5.2.6 Data Retention Profile

The data life-cycle is the expected number of years that the data from a workstream will be stored. These data are all of the remaining non-scratch files from a given workstream. This includes, but is not limited to, source code, input files, run time scripts, generated output, and analysis results.

Table III -- Data Retention Table

	Data Life-Cycle of the data for a given Workstream
WS1	100% for 9 years; 50% as a persistent archive.
WS2	100% for 9 years; 50% as a persistent archive.
WS3	100% for 9 years; 50% as a persistent archive.
WS4	100% the first year; 50% the second year; and 15% past the third year.
WS5	100% the first year; 50% the second year; and 15% past the third year.
WS6	100% the first year; 50% the second year; and 15% past the third year.
WS7	100% for 3 DAYS; 20% the first year; 10% for 3 years; and 5% for 8 years.
WS8	100% for 3 DAYS; 20% the first year; 10% for 3 years; and 5% for 8 years.
WS9	100% for 3 DAYS; 20% the first year; 10% for 3 years; and 5% for 8 years.

The following figure shows the accumulated data retention for WS4 – WS9. This is utilizing the data generation numbers shown in C.5.2.5. WS1 – WS3 cannot be shown without the computational growth being determined.



**Figure 1 - Data Retention for WS4 - WS9**

### C.5.2.7 Automated backup

Automated backup shall be provided for all unique system images on the HPCS; each Home File System (HFS); and the disk-resident inodes, metadata, and files on the HSMS. Software that allows users to restore /home files via a graphical interface is desirable.

For the HFS, a combination of full and incremental backups shall be done to robotically mounted tapes. These backups shall make it possible to restore files to their state on any day during the previous two calendar months. Minimal impact of these backups on the network load is desired.

A history of bimonthly (every 2 months) full backups, for the previous 12 months, shall be produced for shelf storage. A history of annual backups shall be produced until the end of the HPCS system life for shelf storage. At the Government's discretion, shelf storage may be at a remote location. It shall be possible to restore files from these backups until the end of the HPCS system life.

All hardware and storage media used for backup shall be provided by the Contractor.

### **C.5.3 Software Requirements**

#### **C.5.3.1 Resource Management Software**

Efficient operation of the HPCS requires resource management that will facilitate use of the HPCS by NOAA's scientists as well as provide maximum throughput for their workload. The Government desires to implement a variety of charge-back algorithms for monthly processor time or enforce different resource allocations, including disk and tape quotas, for users, groups, or projects on each of the HPCS Subsystems.

The Government requires resource management software that provides:

1. Batch queuing and scheduling (see Section C.5.3.2)
2. Operating system based accounting software, comparable to Unix SVR4 process accounting.
3. System activity monitoring software on the LSC that shows user and system CPU utilization, I/O wait time, and paging activity.
4. Disk quotas on disk space and number of files, comparable to Unix 4.3 BSD quotas.

The Government desires resource management software on the LSC that provides:

- Software partitioning on the LSC

The Government desires the ability to test operating system and application software upgrades in isolation from the interactive and batch production resources on the HPCS.

#### **C.5.3.2 Batch Queuing Software**

The Government requires a Batch Queuing Software package to manage all HPC Subsystem computing resources. The Batch Queuing Software shall be provided with documentation, training, and an established bug-fix process.

To maintain flexibility, the Batch Queuing Software shall support execution hosts with heterogeneous processor architectures and operating systems.

The Batch Queuing Software shall perform the following functions:

- Spool batch job scripts when submitted.
- Tag each job with a user specified string that is at least 45 characters long.
- Run parallel-processed jobs on processors dedicated to the job.
- Create a unique scratch filesystem temporary directory for each job.
- Provide a way for large input files to be pre-staged to the scratch filesystem before the job acquires dedicated processors, and for large output files to be post-staged after the dedicated processors are released.

- Enforce limits on per-process cpu time, job-total cpu time, and elapsed (wallclock) time.
- Track and display job-total cpu time as the job runs.
- Do priority-based job scheduling, allowing higher priority to be assigned to specific jobs, or all jobs owned by a user or group.
- Limit the number of running jobs and monthly cpu hours used by groups and users.
- Collect cluster-global job accounting information for each job, including:
  - the user and group
  - the number of processors dedicated to the job
  - the job-total cpu time
  - the elapsed (wallclock) time
  - the date and time of job completion

The Batch Queuing Software shall provide text-based displays as specified below. It shall be possible to create custom text-based displays, at a minimum by post-processing the display output with text-processing tools such as awk or perl.

- Display all cluster hosts, showing for each host
  - the host up/down status
  - the host load and resource usage
- Display all running jobs, showing for each job
  - the host(s) the job is running on
  - the number of processors dedicated to the job
  - the job-total cpu time
  - the elapsed (wallclock) time
- Display all queued jobs, showing for each job
  - the number of processors requested by the job
  - the date and time the job was submitted
  - the per-process cpu time limit for the job
  - the elapsed (wallclock) time limit for the job
- Display all job accounting records for a specified month

The following Batch Queuing Software features are desired:

- Uniform access to all HPCS computing resources, through use of a single Batch Queuing Software package.
- Meta-scheduling capabilities able to schedule batch jobs on multiple clusters over a wide-area network.
- A project or account feature, with the ability to do priority job scheduling based on projects or accounts, and control user access to the high-priority projects or accounts.
- Priority scheduling of large jobs, keeping smaller, lower priority jobs from starting. In addition, during this draining process, scheduling of small, short jobs, which will complete before the large job starts (backfill scheduling).
- Job scheduling based on consumable resources such as host memory, host disk space, or cluster-global shared disk space.

- Job suspend, job migration, and job checkpointing features that work for FORTRAN90, C, and C++ applications using the proposed parallel programming model.
- Interactive use integrated with the Batch Queuing Software, that provides the batch execution environment to interactive sessions. This will allow batch jobs scripts to be run interactively without change. Accounting of the cpu time used by interactive sessions.

The following Batch Queuing Software features are useful:

- Support for the user-level commands and command options of the POSIX 1003.2d Batch Environment standard (now part of POSIX 1003.1-2001).
- Client software which allows job submission and displays directly from the Government's Linux desktop workstations without interactive login to the HPCS.
- GUI-based displays.

Failover capability for job queuing and scheduling shall be provided.

### **C.5.3.3 Programming Environment Software**

Required application development software for each LSC, development, post-processing and analysis Subsystem of the HPCS programming environment is comprised of:

1. FORTRAN 90/95, C, C++ programming environments, including:
  - a. ANSI standard FORTRAN 90/95, C, and C++ compilers
  - b. macro preprocessors
  - c. source-level debuggers
  - d. performance profilers
  - e. support for 64-bit IEEE reals and integers
  - f. support for reading and writing 32-and 64-bit IEEE floating-point formats in I/O operations
  - g. MPI-1.1, MPI-2 I/O, MPI-2 one-sided communications (for Subsystems supporting parallel environments only)
  - h. The make utility
2. Facilities for source code management, including rcs and cvs.
3. OpenMP is required to the extent that processing units share memory.
4. The Etnus TotalView parallel debugger.

Desired application software for the HPCS programming environments includes:

1. parallelized, optimized numerical libraries on the HPCS
2. optimized (possibly proprietary) I/O libraries.
3. data conversion facilities (for example, endian conversions and conversions of proprietary data formats used in the HPCS). Data conversion facilities will be required if the provided solution requires heterogeneous formats.

#### C.5.3.4 Commercial Off the Shelf Software

The Government requires sufficient software licenses, based upon the user profile, for the following Commercial Off the Shelf (COTS) software. The Government desires that unrestricted licenses are provided.

- Matlab (<http://www.mathworks.com>)
- Mathematica (<http://www.mathematica.com>)
- IDL (<http://www.rsinc.com/idl/index.cfm>)
- S-Plus (<http://www.mathsoft.com>)
- System Performance Monitoring Tools to evaluate load balancing characteristics, system health, and other performance issues.

#### C.5.3.5 Community Supported Software

Required application software for the HPCS programming environments is comprised of:

- GRIB libraries and utilities (<http://www.wmo.ch/web/www/WDM/Guides/Guide-binary-2.html>)
- GNU make (<http://www.gnu.org/software/make/make.html>)
- GNU tar (<http://www.gnu.org/software/tar/>)
- Heirloom cpio (<http://heirloom.sourceforge.net>)
- libxml (<http://xmlsoft.org>)
- netCDF and udunits libraries (<http://www.unidata.ucar.edu>)
- netCDF Operators (<http://nco.sourceforge.net>)
- Prior to the end of the Option Period, the Government expects that its modeling system software will become compliant with the Earth System Modeling Framework (ESMF). More information on ESMF is available at <http://www.esmf.ucar.edu>.
- tcsh (<http://www.tsch.org>)

Required application software for interactive use is comprised of:

- Ferret (<http://ferret.wrc.noaa.gov/Ferret/>)
- Ghostscript (<http://www.cs.wisc.edu/~ghost/>)
- Grace (<http://plasma-gate.weizmann.ac.il/Grace/>)
- GrADS (<http://grads.iges.org/grads/grads.html>)
- ImageMagick (<http://www.imagemagick.org>)
- NCAR graphics (<http://ngwww.ucar.edu>)
- Perl (<http://www.perl.com>)
- Python (<http://www.python.org>)
- Ruby (<http://www.ruby-lang.org>)
- VIS5D (<http://vis5d.sourceforge.net>)

### C.5.4 Networking Requirements

#### C.5.4.1 User Profile

Table IV shows the percentage of HPC Resources allocated for the representative workstream (WSx) according to the locality of scientific users. Scientific users do

not include Computer Operators, System Engineers, System Administrators, and administrative users.

**Table IV - Users of Workstreams by Location (shown in %)**

	WS1	WS2	WS3	WS4	WS5	WS6	WS7	WS8	WS9
Boulder				10%	5%	5%	65%	80%	80%
Washington				80%	80%	80%	10%	10%	
Princeton	96%	96%	96%	5%	10%				
Other	4%	4%	4%	5%	5%	15%	25%	10%	20%

User activities may be grouped into the following four categories:

1. Software and Model Development
2. Data Browsing
3. Model Production
4. Model Analysis

Software and model development includes activities such as code development, debugging and optimization as well as the scientific development of the models. The batch jobs submitted by this user set tend to be short, quick turnaround, but resource intensive (large PE counts, large data sets from archive or long term scratch, etc.) consistent with development of next generation models. Use of editors such as vi and XEMACS, the TotalView visual debugger and visualization tools such as GrADS and nview characterize many of the interactive tools utilized by this user set. The interactive activities require good bandwidth as described below, and “LAN-like” moment to moment quality of service.

Data browsing is an activity ubiquitous to the user environment. Tools such as GrADS and nview are employed to create a graphic representation of spatial and time domain data used as input to the models or created as part of model output. Those engaged in climate workstream activities such as WS1-3 and 5 tend to deal with large data files 1GB to 50GB while weather workstreams such as WS 4,6 and 7-9 with files under 1GB. Such activities require good bandwidth as described below and “LAN-like” moment to moment quality of service. As a further distinction, climate data browsing is generally an “unstructured” activity (i.e. it’s not generally known in advance what data will be accessed) while weather data browsing can be viewed as “structured” (i.e. the desired datasets are known in advance).

Model production is characterized by data staging and model runs generally via the submission of batch scripts. This activity includes post processing such as that for WS1-3, 5 and 8 as well as production of graphics files such as that associated with WS7-9. These production runs can be relatively few, medium to long running, self re-submitting jobs such those characterized by WS1-6 or many short jobs such as WS7-9.

Model analysis is characterized by structured and unstructured data browsing and well as interactive and/or batch jobs to produce verification statistics.

File editing is a ubiquitous activity generally not requiring high bandwidth, but high moment to moment quality of service to meet approximately 25 millisecond response time. Sufficient network bandwidth and quality of service shall be supplied to meet all file editing, data browsing and interactive debugging requirements concurrently. The proposal should clearly indicate how the requirements for data browsing, interactive debugging and file editing at the required frame rates per second and quality of service will be met. Testing of the data browsing, interactive debugging and file editing requirements will proceed throughout the acceptance period with periodic monitoring throughout the contract.

#### **C.5.4.1.1 User Profile for WS7 – WS9**

NOAA has approximately 200 scientific user accounts associated with these workstreams. Approximately half of those users are at NOAA/FSL. Of the approximately 100 users outside NOAA/FSL over 75% are in the Boulder area. The other X% are scattered across the country. There are approximately 40 users engaged in software and model development, 80 users engaged in data browsing roughly 30% of their time, 15 users engaged in model production and 60 users engaged in model analysis. Data browsing tends to be of the structured type on files XXGB using GrADS. GrADS is an X-window based application producing 24bit RGB images with the number of pixels equal to the resolution of the model under study. The data from WS7 (RUC) produces images at a resolution of X1xY1; WS8 (WRF\_SI) produces images at resolution of X2xY2; WS9 (WRF\_Chem) produces images at a resolution of X3xY3. XXX provides an example of the requirement at the high end. A single image produced from XXX data is an RGB image approximately ###MB in size. After initial rendering, an image rate of 12 frames per second implies an uncompressed data rate of XXMB/sec to the display. At time of initial delivery of resources supporting WS7-9, a minimum of NN simultaneous data browsing sessions shall be supported for the Boulder, CO user set. Additionally, a minimum of N simultaneous interactive TotalView debugger sessions shall be supported with LAN-like interactive use for the Boulder, CO user set.

#### **C.5.4.1.2 User Profile for WS4 – WS6**

NOAA has approximately 200 scientific user accounts associated with these workstreams. Approximately 80% of the users reside in the Washington DC metropolitan area. The remainder is spread across the country with some concentration at the NOAA labs in Princeton, Boulder and Miami. There are approximately 50 users engaged in software and model development, 70 users engaged in data browsing roughly 20% of their time, 10 users engaged in model production and 70 users engaged in model analysis. Data browsing tends to be of the structured type on files 1GB or less using GrADS. GrADS is an X window based application producing 24bit RGB images with the number of pixels equal to the resolution of the model under study. The data

from WS4 (WRF-NMM) produces images at a resolution of 501x233; WS5 (GFS) produces images at a resolution of 384X190; WS6 (GSI) produces images at a resolution 768X384. GSI provides an example of the requirement at the high end. Some of the applications represented by workstream 4 anticipate single RGB images approximately 3MB in size. After initial rendering, an image rate of 12 frames per second implies an uncompressed data rate of 36MB/sec to the display. At time of initial delivery of resources supporting WS4-6, a minimum of 30 simultaneous data browsing sessions shall be supported for the Washington, DC user set. Additionally, a minimum of 8 simultaneous interactive TotalView debugger sessions shall be supported with LAN-like interactive use for the Washington, DC user set.

#### **C.5.4.1.3 User Profile for WS1 – WS3**

NOAA has approximately 215 active (used in the last 90 days) scientific user accounts associated with these workstreams. Of the 96% of the Princeton users, approximately 35% are located on Princeton University's campus and are connected to NOAA's laboratory via high-bandwidth connectivity (see Section C.10.4.1). The remaining, approximately, 60% of the active accounts are users at NOAA's lab in Princeton. There are approximately 30 users engaged in software and model development, 80 users engaged in data browsing roughly 25% of their time, 12 users engaged in model production and 80 users engaged in model analysis. Data browsing tends to be of the unstructured type on files 1GB to 100GB using GrADS and ncview. ncview is an X window based application producing 24bit RGB images with the number of pixels equal to the resolution of the model under study. Since the archive is currently mounted to each user workstation, ncview runs as a local application on the workstation, transferring the data from the file sitting on the archive staging disk only once. The data from WS1 (CM2-ESM) produces images at resolutions of 144X90 (atmosphere and land) and 360X200 (ice and ocean); WS2 (CM2-HR) produces images at resolutions of 288X180 (atmosphere and land) and 1080X840 (ice and ocean); WS3 (HIMF) produces images at resolution 2160X680. HIMF provides an example of the requirement at the high end. At the benchmark resolution, a HIMF global field is approximately 1.47 million – 64-bit real values; the netCDF data file containing the field may be 10-100GB. A single image produced from this field by ncview is read, transferred and rendered in about 1.5 seconds producing an RGB image approximately 4.4MB in size. After initial rendering, an image rate of 15 frames per second implies an uncompressed data rate of 66MB/sec to the display. At time of initial delivery of resources supporting WS1-3, a minimum of 30 simultaneous data browsing sessions shall be supported for the Princeton, NJ user set. Additionally, a minimum of 10 simultaneous interactive TotalView debugger sessions shall be supported with LAN-like interactive use for the Princeton, NJ user set.

**C.5.4.2 Wide Area Network (WAN) Component**

NOAA has existing WAN capabilities as described in Appendix A. The vendor is responsible for providing any additional connectivity required by the location of a provided Subsystem. For example this requirement may be driven by the data generation on the LSC as described in the Table II in Section C.5.2.5 or between the post processing and analysis component and the data and archiving component.

NOAA's Enterprise Network Target Architecture (ENTA) is currently being developed. The resulting NOAA-wide network infrastructure could be referred to as NOAAnet. Upon completion of this architecture, the Government may, at its option and at any time during this contract, require the Contractor to alter the WAN connection to comply with the NOAA ENTA. Such a change will be negotiated.

The Contractor is encouraged to be innovative in their Wide Area Network solution in order to optimize total system performance. The Government desires to balance the Wide Area Network solution that is provided with the impact it will have on the interactive experience of remotely located users.

**C.5.4.3 High bandwidth connectivity to model and observation data**

Near real-time observation and model data are required for use by given workstreams. The model data will be created and will be available from the NOAA Operational HPCS. The satellite data will be available from both the Operational HPCS and the NESDIS CEMSCS computer that is located in Suitland, MD. Table V shows the projected input data quantities required by each workstream. All model and observational data will originate from the NOAA operational data stream. Note that workstreams 4 and 6 require all of the data used by workstreams 7-9. More specifically, the data required for workstreams 7, 8 and 9 are a subset of that required for workstreams 4 and 6. Thus, if multiple workstreams can access the observation and model data sets from workstream 4, only a single transfer of the information would be necessary for those workstreams. There may be additional needs for real-time data acquisition from other sites, but they would be funded externally to this contract.

**Table V- Data Ingest (TB/day)**

	Required by WS7, WS8, and WS9	Required by WS4 and WS6 in addition to that required by WS7, WS8, and WS9
FY2006	1	
FY2007	2	1
FY2008	2	1
FY2009	4	2
FY2010	4	2
FY2011	8	4
FY2012	8	4
FY2013	12	12
FY2014	12	12

Approximately half of the data are model data. The other data is composed of operational and experimental satellite data, which are expected to be used in data assimilation systems. The expected data sources are the current satellites still in use at the beginning of the contract possibly including METEOSAT, GOES-10-12 and NOAA-14-18 (both imagers and sounders) as well as NASA's EOS-Aqua and EOS-Terra satellites (MODIS on both and AIRS/AMSU-A on Aqua), SSM/I (imager data) and SSM/IS (F16) on the DMSP satellites. New polar orbiting satellite data are expected to come from NOAA-N', NPP (sounder and imager in 2007), SSMIS on F17 (2005), F18 (2007), F19 (2008), F20 (2010), METOP (sounders and imager in 2005), NPOESS (sounders and imagers) C1 (2010), C2 (2011) and C3 (2013). For geostationary satellites, new sounder and imager satellite data will include GOES-N (2005), -O (2007), -P (2008) and -R (2012). All launch dates are approximate, and some need for transfer of simulated data streams of the same size as the actual data sets may occur up to a year prior to the launch dates.

Provision of the data to the NOAA R&D HPCS does not require the same timeliness as for operations. A delay of the data transfer of up to 6 hours (beyond that from the operations) is acceptable. At the discretion of the Government, further delays of the data can be counted as machine down time. Also, it is required that no significant gaps in the data sets occur because of communications and any delayed data should be communicated as soon as possible.

Currently available bandwidth should make the necessary data available within the Washington D.C. metro area and to the Skaggs building in Boulder. It is uncertain if the necessary bandwidth will be available to GFDL in Princeton. The Contractor is required to provide any necessary bandwidth, above and beyond the bandwidth provided by the Government, for the communication of this data.

### **C.5.5 IT Security Requirements**

To assure an adequate level of protection for in-house or commercially maintained IT systems, NOAA maintains all systems consistent with government-wide laws and regulations. The Office of Management and Budget (OMB) Circular A-130 requires all federal agencies to plan for the security of all IT systems throughout their life cycle. OMB Circular A-130 requires agencies to implement and maintain a program to assure that adequate security is provided for all agency information collected, processed, transmitted, stored, or disseminated in general support systems and major application. It further directs each agency to implement policies, standards, and procedures that are consistent with government-wide policies, standards, and procedures issued by the Office of Management and Budget (OMB), the Department of Commerce (DOC), the General Services Administration (GSA) and the Office of Personnel Management (OPM). OMB Circular A-130 also establishes a minimum set of controls to be included in Federal automated information security programs and assigns Federal agency responsibilities for the security of automated information.

At a minimum, agency programs shall include the following controls in their general support systems and major applications:

- Assignment of Responsibility for Security.
- System Security Plan consistent with guidance issued by the National Institute of Standards and Technology (NIST) to include:
  - Rules of Behavior/Application Rules.
  - Training.
  - Personnel Controls.
  - Incident Response Capability.
  - Continuity of Support/Contingency Planning.
  - Technical Security/Controls.
  - System Interconnection/Information Sharing.
  - Public Access Controls as required.
- Scheduled Review of Security Controls commensurate with the acceptable level of risk to the system.
- Authorize Processing (Certification and Accreditation).

The Federal Information Security Management Act of 2002 (FISMA), (HR 2548 E-Government Act, TITLE III – INFORMATION SECURITY, SEC. 301. INFORMATION SECURITY), addresses the program management and evaluation aspects of IT security.

Additionally, NOAA follows the US DOC Manual of Security Policies and Procedures (April 4, 2003), the US DOC IT Security Program Policy and Minimum Implementation Standards (July 28, 2004), as well as NOAA OCIO Policies and Directives.

More information can be found at:

- <http://www.cio.noaa.gov> (NOAA CIO web-site)

- <https://www.csp.noaa.gov> (NOAA IT Security Office web-site)
- <http://www.osec.doc.gov/cio/oipr/ITSec/ITSECDOC1.HTM> (DOC IT Security Program Office web-site)
- <http://csrc.nist.gov/policies/> (NIST Federal Requirements)
- <http://csrc.nist.gov/publications/nistpubs/index.html> (NIST Guidelines)

In order to maintain a consistent security posture across potentially multiple sites, the Contractor shall be required to implement DOC/NOAA equivalent Patch Management, Change Management, and Perimeter Protection procedures for each processing site. These procedures shall be documented and maintained as part of the required NIST SP 800-37 Certification and Accreditation package. Although the R&D HPCS will be managed as a single system regardless of physical location, it is expected that multiple sites must rely on public Internet infrastructure for inter-site communication. Special care must be taken to secure inter-site communication over an untrusted network, and mitigate the risks associated with the inherent trust relationships required for full R&D HPCS integration.

The Contractor shall provide system documentation, which contains descriptions of the system hardware, software, policies, standards, procedures, and approvals related to the system life cycle and a formalized process of the system's security controls. Per NIST SP800-37, the system must be recertified after undergoing a major configuration change, therefore the mid-life upgrade cycle will require recertification of the C&A package.

In addition, all solutions shall comply with all DOC and NOAA security policies and industry best practices will be used to increase the security posture of the system beyond the DOC and NOAA policy requirements. NIST guidelines/practices are preferred industry best practices. For planning purposes, the R&D HPCS system will be designated as NOAA Mission Critical, with a FIPS 199 security categorization of {(confidentiality,low), (integrity,moderate),(availability,moderate)}. The resulting moderate-impact level designation may then be used to select recommended controls from NIST SP-53 (“Recommended Security Controls for Federal Information Systems”).

### **C.5.6 Facilities Requirements**

The Government does not mandate a facility solution. Rather, the Government requires the Contractor to provide a facility proposal for housing the proposed equipment to meet the Government’s requirements over the life of the contract. The facility proposal must provide flexibility and provide for disaster recovery. In the proposal, Contractors may propose to use any of the following to meet the Government’s needs: any of the available Government facilities (as described in Appendix B), either as is or with modifications; Contractor-provided facilities; or a combination of Government facilities and Contractor facilities.

In their facilities plan, Contractors may use the projected facilities resources offered at three NOAA locations, BLDR-1, BLDR-2, and PRTN, which are described briefly below and in more detail in Appendix B indicates the government-supported resources that are projected to be available, beginning at the times indicated and are provided under the assumptions indicated in that section. If Contractors require additional facility resources at these locations beyond what is indicated, they must provide a modification plan indicating how they will obtain these resources. The modification plan must be in accordance with site restrictions.

Contractor use of the government-supported resources implies the following:

The Government is responsible for providing the indicated facilities resources. Therefore, any lost system time caused by environmental outages or facility failures will be recorded as null time for availability calculations.

The Contractor is responsible for any facility modifications (e.g., new electrical equipment, plumbing, additional air conditioning equipment) needed to support the proposed systems to be installed. A description of these modifications must be included as part of the Contractor's facility plan.

The following are brief descriptions of the Government facilities:

Two of these facilities, designated as BLDR-1 and BLDR-2, are located in the David Skaggs Research Center (DSRC), 325 Broadway, Boulder, CO, and the third, designated as PRTN, is located in the Geophysical Fluid Dynamics Laboratory (GFDL) building complex, referred to as the "Princeton Complex" in Section C.8 at 201 Forrestal Road, Princeton University Forrestal Campus, Princeton, NJ.

BLDR-1 currently houses NOAA's JET computer system, which is operated by OAR's Forecast Systems Laboratory under a contract with HPTi. The BLDR-1 facility will not be available until October 2006. BLDR-2 is a facility currently under construction at DSRC and will be available for use in October 2005. PRTN currently houses NOAA's HPCS operated by OAR's Geophysical Fluid Dynamics Laboratory under a contract with Raytheon Company

Appendix B provides descriptions of these government-furnished facilities that are designed to support high-performance computer systems. These sites are available for possible use by Contractors to house systems under the NOAA HPC R&D contract, subject to the availability of facility resources as projected in Section C.8.11 under the assumptions provided therein.

The following table summarizes projected raised floor space by location and earliest date of availability:

**Table VI – Raised floor space**

Projected Resources	Site	October 2005	October 2006
---------------------	------	--------------	--------------

Raised Floor Space (sq.ft.)	BLDR-1	0	2250
	BLDR-2	1500	1500
	PRTN	2700	5500

The Government will offer site visits to Offerors on the following schedule:

BLDR-1 and BLDR-2	January 18, 2005
PRTN	January 20, 2005

During the site visit, Offerors may obtain a set of blueprints for that facility once an authorized representative signs a non-disclosure agreement and agrees to destroy all blueprints at the end of the solicitation. Note in particular that Offerors requesting building diagrams or related documents of federal government facilities (specifically all documents relating to DSRC) shall be required to have an authorized representative execute a non-disclosure form as required under GSA PBA-3490.1.

The Contractor may propose to use facilities other than the Government-provided facilities as described above. If the Contractor chooses this approach, the provided facilities will have the following conditions:

- Contractors are responsible for providing all facilities resources, including floor space, utilities, facilities maintenance, janitorial services, etc. Therefore, any lost system time caused by environmental outages (such as loss of power, cooling, etc.) or facility failures will be recorded as downtime for availability calculations.
- Contractors may use any Unrestricted GFE provided in Appendix C, but are responsible for all shipping costs, both to the Contractor site and back to a Government directed site at the end of the contract and loss of use charges during shipment.
- The Contractor may assume credits against contract costs according to the following terms, if he does not choose to use the offered facilities:
  - BLDR-1 and/or BLDR-2: No credits will be offered for not using these facilities
  - PRTN: If the Contractor choose not to use the PRTN facility, the Government will provide credits within the contract expenditures to cover the cost of power at the proposed facility of up to 6,500,000 kilowatt hours
- Designated Government personnel must have suitable access to the Contractor facility on an as-needed basis, subject to reasonable vendor restrictions.
- The Government will maintain a description of the facilities housing the R&D HPCS. Contractor provided facilities shall require that the Contractor maintains the description. Government provided facilities shall be maintained with the assistance of the Contractor. The facilities description must include:
  - 1) A copy of the site operating plan
  - 2) A copy of the security procedures
  - 3) A guarantee of the maximum number of facility incidents and annual outage time

- 4) A statement documenting any single points of failure in computer conditioning
  - 5) A description of how NOAA remote computer operators will be informed of deteriorating facility conditions such as rising room temperatures or an air handler failure
  - 6) A copy of the disaster recovery plan
  - 7) One-line (logic) diagrams of the electrical service and cooling service
  - 8) An energy density (watts per square foot) projection plotted over the contract life
  - 9) A spreadsheet listing the type and age of facility equipment to be used. Examples are: UPS systems and power conditioners, chillers, heat exchangers, air handlers.
  - 10) A copy of the contract statement of work for any commercial facility management company used, or the equivalent if performed in-house. Preventative Maintenance schedules, proactive inspections, and quality assurance methods are examples.
  - 11) A brief (2-3 paragraphs) description of the procedures used to acquire off-site emergency service including minimum response times and escalation procedures
  - 12) A statement (one paragraph) as to how coverage and services are made available after-hours and on holidays
  - 13) A statement (paragraph) projecting the minimum UPS power protection period (survival time) when utility power fails. The worst-case instance covers full functionality of the system(s), and over the entire period of performance (i.e., should load vary during the contract, or should battery performance decline due to age).
  - 14) A brief description of the fire protection systems and certification standards
  - 15) A list of design standards complied with by the facility architect (see above)
  - 16) A brief description of any automated facility controls such as computer-managed failover systems
  - 17) A brief description of quality assurance control processes for defect corrections
  - 18) A brief description of the tools used by management to track performance of service level commitments and to manage the facility (Examples: CAD programs, equipment inventory management software, capacity planning and tracking, etc.)
  - 19) A bio (curriculum vitae) of the facility manager's experience and training
  - 20) A description of facility alterations and changes to be made to the offered space if the Offeror is successful
  - 21) Cost data as described above
- The Government reserves the right to conduct site visits of all Contractor provided facilities. During the site visit, the Government may inspect:
    - 1) Mechanical rooms

- 2) Raised floor plenums; drainage; cable tracks, labeling, and management
- 3) Network cabling protection and redundancy provisions
- 4) Power and cooling distribution and control systems
- 5) Logs of equipment failures, corrective actions taken and maintenance results
- 6) Preventive maintenance schedules
- 7) Emergency plans, and results of drills
- 8) Testing procedures and schedules
- 9) Training provided to facility managers, maintenance and security personnel
- 10) Safety and fire protection equipment and operation
- 11) Any other relevant materials that will enable the government to assess the reliability and safety of the facility

**C.5.6.1      *Procedures for government-designated personnel to access contracted facilities***

If the contractor proposes the use of non-Government facilities, appropriate procedures that allow government-designated personnel to access these facilities must be provided.

**C.5.6.2      *Policies for providing workstations, phone service, office equipment to on-site contract personnel***

The following are policies that will be in force if the Contractor chooses to use the designated facilities.

**BLDR-1 & BLDR-2**

The Government provides phones and LAN connections via Standard Information Outlets (SIO) only. All workstations and office/lab furniture will be supplied by the Contractor.

**PRTN**

The Government requires the vendor to provide phone service, workstations, and office furniture for Contractor Support Personnel. However, Contractor-supplied software analysts, security specialists, and network support personnel will be provided workstations, office furniture, and phone service to the extent that they are integrated into the laboratory's normal office infrastructure.

**C.5.6.3      *Procedures for obtaining construction permits***

The following are procedures that must be followed if the Contractor chooses to use the designated facilities.

**BLDR-1 & BLDR-2**

Construction permits are not required because DSRC is a federal building.

**PRTN**

Under the current contract, any construction to be done requires the Contractor to obtain construction licenses, permits, and approvals for this work from Plainsboro Township and Princeton University, as required.

**C.5.6.4 Costs for facilities modifications in GSA properties**

GSA Regulations require that GSA Contractors carry out all modifications to GSA buildings. Therefore, any facilities modifications will require that funds be held back from the contract for transfer to GSA for any changes to be made to GSA properties.

**BLDR-1 & BLDR-2**

Close coordination with the DSRC landlord, U.S. General Services Administration (GSA), will be required if any building modifications are required. Costs will include normal GSA overhead charges plus project management oversight. All modifications must meet standard building codes as dictated by GSA. Modifications can include small projects like power receptacle changes, to large projects like chiller plant upgrades. Normal government contracting procedures can be expected for any modifications.

**PRTN**

Not applicable

**C.6 Reliability and Availability Requirements****C.6.1 Reliability and Availability**

The HPCS shall continue NOAA's historically high utilization of its computing resources. Reliability, availability, and Contractor support are considered fundamental aspects of the HPCS.

**C.6.1.1 Reliability**

Downtime will be used in the determination of the actual System Life Throughput. Downtime is that period of time when all of an HPC Subsystem's workload cannot be accomplished due to a malfunction in the Contractor-maintained equipment or software, or when the HPCS or one of its components is released to the Contractor for maintenance. Periods of Remedial and Preventive Maintenance count as downtime. Null time is that period of time when the workload cannot be accomplished due to environmental failure, such as loss of electric power or cooling, or recovery from environmental failure. Null time will not be counted as downtime. Downtime for each HPCS component is based on the fraction of the resources available for that component's workload. It is arrived at through consultation between the Government and the Contractor, and determined ultimately by the Government. Downtime is accumulated on the HPCS if the Government is not able to execute its workload when access to the scientific data via the storage or archiving components is unavailable.

A component's downtime shall commence at the time the Government makes a bona fide attempt to contact the Contractor at the designated point of contact (see Section C.7.1). At this time, the Government will begin a log of the problem which will be completed and signed by both the Government and the Contractor when the problem is resolved. Information to be entered into the log will be determined by the Government.

A component's downtime shall exclude any time in which the Government denies the Contractor maintenance personnel access to the malfunctioning hardware and/or software. (Routine procedures for entry to a facility shall not be construed as denial of access.)

A component's downtime shall end when the computer is returned to the Government in operable condition as determined by the Government, ready to perform all of the workload.

Preventive Maintenance (PM) is to be completed at times determined by the Government.

The testing and installation of every major operating system release installed at the request of the Government and one (1) minor operating system release installed at the request of the Government during any annual period will count as downtime. Preparation for and execution of post-upgrade LTDs, including any benchmark development by the Contractor, associated with the agreed-upon upgrades proposed at contract award, will count as downtime.

#### **C.6.1.2 Availability**

All components of the HPCS must perform as an integrated system to provide the Government with at least 96% availability. These components include not only those typically associated with a high performance computing system but also any WAN component provided as part of the HPCS. Additionally, the Government requires 99% availability for access to its scientific data. Scientific data archives may be on the Long-Term Scratch component or the HSMS component, depending on the overall architecture associated with a given workstream. The Government's data are critical assets associated with the HPCS. With the large amount of data associated with the HPCS, the Government requires the highest levels of data integrity.

Although availability will be monitored daily, it will be measured on a monthly basis. The Contractor shall establish and maintain accurate system monitoring and/or logs to support system availability measurements. Monitoring tools must be made available to the Government. Throughput shall be determined by relating system availability to the workstreams normally run on that system. For example, if Subsystem A typically runs workstreams 1 – 3, workstreams 4 – 9 will not be used to determine the amount of "lost" throughput, if any.

Availability shall be determined by computing the ratio of total computation processor hours available for execution of R&D jobs to the total computation processor hours each month, with regards to Null Time. The time a computation processor is available for execution is determined by subtracting processor downtime from wall clock time. Spare processors can be included in the computation pool to reduce downtime. However, downtime accumulates until the spare processors are made available for job execution. Accumulated computational cycles (in CPU-hours) that are lost when jobs are lost due to component failure or component reboot will not count toward the system-life throughput calculation. If the accounting software cannot report the accumulated computational cycles for each active job at the time of failure, it will be assumed that four (4) CPU-hours were lost for all processor(s) on the failed system.

System Life Throughput (SLT), for a given workstream  $i$ , can be calculated by the following equation:

$$SLT_i = \sum_j \frac{T_{i,j} A_{i,j}}{B_{i,j}}$$

$SLT$  = System Life Throughput  
 $T$  = total wall-clock Time during system configuration  $j$   
 $A$  = proposed Availability  
 $B$  = workstream Benchmark time  
 $i$  = WS number  
 $j$  = system configuration period

An example of System Life Throughput for a sample workstream can be demonstrated for FY2001-FY2003. This calculation is based on actual calendar days beginning October 1, 2000. For example, if the proposed LSC throughput benchmark execution time is 10,800 seconds initially, and is upgraded to 7,200 seconds on October 1, 2001, and a 98% availability is proposed through the three-year period, the system life throughput SLT is given below (Note: there are 86,400 seconds in a day).

$$SLT_{sample} = \frac{((365 \frac{days}{year})(86400 \frac{s}{day}))(0.98)}{10,800 \frac{s}{benchmark}} + \frac{(2(365 \frac{days}{year})(86400 \frac{s}{day}))(0.98)}{7200 \frac{s}{benchmark}}$$

$$SLT_{sample} = 11446.4 benchmarks$$

Proposed throughput benchmark performance levels will be combined with the proposed availability level (98% is shown in the example above) to determine a measure of overall proposed System Life Throughput for the workstream. The actual throughput will be measured on a daily basis by monitoring the availability of the components associated with the given workstream. All performance levels must be met for each measurement of actual throughput regardless of past delivery of suites.

Although system throughput shall be monitored daily, the accumulated System Life Throughput of the workstream shall be calculated monthly. Shortfalls shall be made up with new equipment brought in at no additional cost. Using the demonstrated benchmark performance on the upgraded HPCS, the Government will calculate how long the upgrade shall stay in place to compensate for the shortfall in throughput. It is the Government's goal to meet the total System Life Throughput by the end of the contract in a manner that does not require frequent disruptions, front-loading of cycles, or back-loading of cycles.

At the option of the Government, shortfalls in throughput on the HPCS due to downtime shall cause downtime credits to accrue. These downtime credits shall be in lieu of bringing in new equipment. Downtime credits shall also accrue on all Subsystem-wide available file systems.

To better reflect NOAA's computational needs over time, changes in the HPCS workstream benchmarks shall be made by mutual agreement between the Government and the Contractor through the life of the HPCS.

The Government requires a credible method of maintaining system availability. The Government expects that fail-over and high reliability components will be used. The Government requires sufficient power during environmental failure to gracefully shut down all components of the HPCS. Further, the Government requires adequate power conditioning to insulate the System from power spikes and sags. It is expected that uninterruptible power supplies (UPSs) and power distribution units (PDUs) will be needed for all components of the HPCS to meet these requirements.

## ***C.7 Support Services Requirements***

### **C.7.1 Support**

The Contractor shall provide the Government with a designated point of contact to request maintenance. The Contractor shall maintain escalation procedures that allow the Government round-the-clock telephone contact with knowledgeable Contractor staff should the designated point of contact be unavailable.

For each provided HPC Subsystem, the Government requires comprehensive support in order to meet the 96% availability requirement. Support professionals can consist of systems analysts, hardware engineers, and applications analysts. Contractor system analysts shall work closely with Government system administrators. Government system administrators shall have root access to the HPCS computing platforms. The Government requires a senior applications analyst to be located at the primary user sites for WS1 – WS9. Senior Application Analysts shall be available, in person, during business hours. The primary user sites are described in the User Profile Section, C.5.4.1. NOAA expects to provide offices for the senior analysts at its facilities (see Section C.5.6.2 for more information). It is

anticipated that additional senior application analysts may be required during transitions to new technologies (see Section C.9.5).

The Government requires an itemized list of all Contractor-supplied hardware and software items, and documentation of these items, in printable electronic form.

### **C.7.2 Training**

Training shall be provided for NOAA computer specialists and operators in:

- system administration and tuning
- hardware operation and system overview

Training shall be provided for a large number of NOAA applications programmers in:

- application and shell programming
- programming languages and tools
- HSMS software user interface
- optimization

The Contractor shall provide the Government with a list of additional potential training topics. Training must be collocated with users to the greatest extent possible. On-line training with access to experts will be considered.

The Government desires to begin training when early access to systems similar to those proposed for the HPCS is granted. See Section C.5.1.2 for more information.

## **C.8 Project Plan Requirements**

### **C.8.1 Project Management**

The contractor shall be responsible for all functions related to managing the NOAA HPCS R&D system as a single project. NOAA will retain responsibility for project oversight and directing project activities (such as resource allocation) that directly impact the Government. The contractor shall establish communication avenues to keep NOAA managers apprised of daily status and alerts. A formal project management review process shall also be established as well as described below.

Reference materials for standard Government Information Technology project management are available on the Internet, for example:

<http://cio.doe.gov/ITReform/sqse/publications.htm#Checklists>

The contractor shall appoint a Project Manager and assistants as required. This management team will be responsible for meeting NOAA's technical and business requirements (including security), project planning (cost, schedule quality), resource mapping, identifying dependencies and support issues and managing subcontractors. The Project Manager shall hold an Annual Meeting for NOAA management in Silver Spring, MD to discuss project status and plans for the coming year.

### **C.8.2 Transition to “One NOAA”**

Transition to “One NOAA” NOAA requires increased integration of its HPC resources over time to provide additional flexibility and robustness. Elements of the long-term vision include:

- 1) the ability of users to access any computational platform in order to:
  - i. achieve code interoperability with good performance
  - ii. provide continuity of operations
  - iii. respond to changing programmatic requirements
  - iv. Provide similar programming environment and tools across platforms to minimize user re-training.
- 2) the ability of users to access data from any platform of the HPCS and to migrate data from any existing archive to future archives
- 3) an ongoing collaboration with the Contractor to produce additional business processes that promote integration

Features of the NOAA R&D HPCS that promote this vision may include a single security profile, including single sign-on, and a single user interface into all HPC resources, which may include a single batch scheduler or a metascheduler. NOAA recognizes that not all elements of the vision may be available on the system provided during the term of the base contract, or may compromise performance unacceptably. In this case, NOAA may choose to forego implementation of one or more of the elements of this vision.

NOAA considers the risk of transition to the full suite of these capabilities at day one of the contract to be very high due to both the immaturity of the underlying technology for distributed operations and the significant business process changes required of users and institutions.

### **C.8.3 Documentation**

The Contractor shall provide adequate documentation for maintenance of all Project Management and the Configuration Management system. This includes, but is not limited to, baseline configuration, implementation methodology, a formal project plan, and documentation of administrative tasks, etc....

### **C.8.4 Contract Transition Plans**

Three months before any significant hardware changes, the Government requires a Transition Plan. After the change, the Government requires post-mortem review.

### **C.8.5 Configuration Management Plan**

The Contractor shall establish a Configuration Management (CM) process including a Configuration Management Control Board to include both contractor and NOAA staff. The Government desires that the Control Board and Configuration Management process integrate into existing Government Control Boards and processes where applicable. Configuration management plans (at the minimum for network, hardware and software (see Section C.4.2)) shall be developed by the contractor and presented to the CM Board for approval. Not less than monthly CM

meetings shall be held to review system status and plans for the future. The contractor shall be responsible for preparing an Annual Configuration Management Plan for presentation to NOAA management in Silver Spring, MD. The Configuration Management Control Board will be responsible for oversight of benchmark workstreams, modifications to the benchmarks and for scheduling system time in support of benchmarking.

## **C.9 Contract Options**

### **C.9.1 Option Period**

The Base Period of the contract is FY2006-FY2009. The Government requires an option to extend the contract for another four years from FY2010-FY2013. Requirements are listed in Section C.

### **C.9.2 One-year extension of Base Period**

After the four-year Base Period, for FY2010, the Government requires the option to extend the lease for operations, maintenance and the equipment for one additional year in quarterly increments. The Government will find it useful for funding to be half of the FY2009 funding level (see Section C.4.3).

### **C.9.3 One-year extension of Option Period**

After the four-year Option Period, for FY2014, the Government requires the option to extend the lease for operations, maintenance and the equipment for one additional year in quarterly increments. The Government will find it useful for funding to be half of the FY2013 funding level (see Section C.4.3).

### **C.9.4 Additional R&D HPCS Augmentations**

Over the life of the contract there may arise situations that will require the R&D HPCS to be augmented. One such situation might result if NOAA identifies a new requirement for HPC that did not exist at the time of contract award. A second situation might result if NOAA were to enter into an inter-agency agreement with another Federal agency to supply that agency with computational resources. In either situation the Government will request the Contractor to provide a proposal to meet any such requirement.

### **C.9.5 Engineering Support**

The Government requires additional expert-level engineering to address impending needs for a given workstream. Due to off-site meetings or presentations, some travel may be required as needed.

#### **C.9.5.1 Applications Analyst**

NOAA may require senior applications analysts for porting, tuning, optimizing, and developing scientific applications in numerical weather prediction or climate prediction. The Applications Analyst will work with NOAA Application engineers,

and Scientists. The Analyst is required to have direct access to vendor compiler development teams.

Engineer Qualifications:

Seven years of HPC experience in porting, tuning, optimizing, and developing scientific applications in numerical weather prediction or climate prediction. Must have expertise in FORTRAN 77/90/95/coarray. May need expertise in C and C++. The Government desires an individual with direct experience in Physical Sciences that relate to Weather or Climate Research.

**C.9.5.2 Systems/Network/Security Engineer**

NOAA may require professional engineering services in support of its HPC facilities. These services may be for System, Network, or Security Engineering. Activities may include, by are not limited to the following:

- Install/remove vendor supplied OS and systems software (compilers, batch system, etc) upgrades.
- Manage upgrades for COTS or Community Supported software supplied via the contract.
- Use systems utilities to tune performance.
- Provide utilization and performance information to the extent the systems allow.
- Customize configurations of software and hardware.
- Collaborate with Contractors and Government staff.
- Comprehensively documenting plans and performed activities.
- Diagnose technical problems. Including distinguishing between hardware, software, and user errors.
- Work at an application level in the event job scheduling, interpretive utilities (such as scripting languages), compilers and runtime libraries are suspect.
- Support end users.
- Systems Analysis.
- Evaluate current methodologies and offer expert level advice on alternatives or target solutions.
- Comprehend and apply security upgrades. Collaborative work with Government IT security staff is essential.

Engineer Qualifications:

Education typically includes a minimum of a Bachelor's Degree in Computer Science, Engineering, or a closely related field, however appropriate experience can be substituted. Typically at least five years of senior-level experience is required. Has successfully completed a 40 hour security training course. Has successfully completed a background investigation.

The Government desires the following Additional qualifications:

- As applicable, industry standard certifications. Such as SANS GIAC, CCIE, etc.
- Is an active member of USENIX or SAGE and has recently been to a LISA (Large Installation System Administration) Conference.

### **C.9.5.3 Facilities Engineer**

NOAA may require professional engineering services in support of its HPC facilities. These services may be architectural, electrical, mechanical or civil engineering services. Activities may include, but are not limited to the following:

- Facility Engineering Consulting Services
- General Engineering Studies
- Specific Engineering Studies
- Design Services
- Design-Build Services

#### Engineer Qualifications:

All services will need to be performed by engineers who meet the following minimum requirements:

- 10 years experience in High Performance Computer (HPC) facility engineering.
- Past history of successful HPC facility engineering services.
- Specific HPC power, cooling and infrastructure experience.
- Membership in good standing in appropriate engineering trade organizations.
- Past performance in Federal Government contracting.

#### The Government desires the following Additional qualifications:

- Professional Engineering (PE) services.
- Engineering Stamps.
- Plan draft, review and approval.
- Drawing update and certification.

## **C.10 Appendix A – Details of current NOAA R&D HPCS**

### **C.10.1 Large-Scale Computing (LSC)**

#### **C.10.1.1 Boulder, CO**

NOAA's Large Scale Computing located in Boulder consists of five logical clusters based on Intel 2.2GHz Xeon systems interconnected with MyriNet. Each node has two CPUs and 1 GB of memory. Two clusters contain 62 nodes each, one cluster contains 128 nodes, and two clusters contain 256 nodes. There are four front-ends with 4 GB each. There is also a test bed containing 12 dual-processor Opteron systems and 12 dual-processor Itanium systems.

#### **C.10.1.2 Washington, DC**

NOAA's Large Scale Computing located in Washington, DC consists of 40 IBM Power4 1.3 GHz Regatta H compute nodes, with a total of 1280 processors. Each node has 8 Logical Partitions (LPARs) with 4 processors per LPAR. Each processor has 1 GB of memory. The interconnect is an IBM SP Switch-2, which provides for multiple plain support for the switching fabric. Each LPAR has 2 SP Switch-2 PCI adapters.

The storage and I/O system is integrated with the switch and contains 4 additional dedicated Regatta H nodes each with 32 processors, configured in 4 LPARs with 8 processors per node. Each processor has 2 GB of memory. Each LPAR has two IBM SP Switch-2 adapters.

#### **C.10.1.3 Princeton, NJ**

NOAA's Large Scale Computing located in Princeton is designed for the batch processing of computationally intensive jobs. The LSC consists of 11 nodes with 3040 processors. Fast scratch disk is available from each node. The fast scratch disk has eight 1 Gb FC per host and is in a RAID 5+1 configuration.

There are eight SGI Origin 3000s running Irix 6.5.19f and consisting of:

2 hosts with	512 x 600 Mhz MIPS, 512 GB memory, 2.6 TB fast-scratch disk
5 hosts with	256 x 600 Mhz MIPS, 256 GB memory, 0.6 TB fast-scratch disk
1 host with	128 x 600 Mhz MIPS, 128 GB memory, 0.6 TB fast-scratch disk

There are three SGI Altix 3700s running Red Hat Enterprise AS with SGI Propack and consisting of:

2 hosts with	256 x 1.5 Ghz Intel, 512 GB memory, 1.4 TB fast-scratch disk
1 host with	96 x 1.5 Ghz Intel, 192 GB memory, 2 TB fast-scratch disk

In April 2005, NOAA will receive a performance increment of at least 1.8X above the Origin portion of the Large Scale Computing component.

## **C.10.2 Post-Processing and Analysis**

### **C.10.2.1 Boulder, CO**

There are four visualization servers based on Intel 1.7 GHz Xeon systems. They have dual ethernet connectivity (to the internal cluster network and to the internet via NOAA/FSL's LAN). There are also four cron servers based on 2.2 GHz Xeon systems.

### **C.10.2.2 Washington, DC**

NOAA's HPCS system located at Washington DC contains 16 LPARs that are dedicated to single threaded batch jobs (mostly pre- and post-processing) and for interactive access to the parallel compute environment, which consists of the remainder of this HPCS. Additional post-processing and visualization is performed at several branch servers at NCEP, which are typically Origin 300 or Altix servers.

### **C.10.2.3 Princeton, NJ**

NOAA's Analysis system located at Princeton has two nodes for analysis, post-processing and development work. The Analysis system is comprised of two SGI Origin 3900 systems and one SGI Onyx 3000 with visualization tools.

2 hosts with	96 x 600 Mhz MIPS, 96 GB memory, 3.9 TB fast-scratch disk
1 host with	4 x 400 Mhz MIPS, 4 GB memory, 0.2 TB scratch disk

The fast scratch disk is comprised of 48 x 1 Gb FC per host that connects to a XFS filesystem in a RAID 5+1 configuration.

The analysis computers are available for interactive use at all time. Parallel applications on up to 16 CPUs may be run for up to 8 hours. For applications that use more than 16 CPUs, use of the Large-Scale Computing is recommended, not required.

## **C.10.3 Storage and Archiving**

### **C.10.3.1 Home File Systems (HFS)**

#### **C.10.3.1.1 Boulder, CO**

There are two Data Direct Networks S2A 6000s each has 8 x 1 Gb/s FibreChannel ports. One has 9.2 TB of useable space and the other has 6.7 TB of useable space (this space excludes parity stripes). Files systems are served by a number of Linux-based NFS servers with read/write data rates of 40-50 MBytes/second each. There are six Dell 2650s with dual 1 GHz Pentium IIIs, one Dell 2550 with dual 933MHz Pentium IIIs, three dual-processor 2.2GHz Xeon systems, and one dual-processor 2.6GHz Xeon system. All of the systems have 2 GBytes memory, one Gigabit ethernet interface and one 1 Gb/s FiberChannel HBA. These file systems provide long

term (but not backed up) data storage. These file systems also provide the /home file system for the clusters.

There is a 1.6TByte scratch file system based upon PVFS that utilizes commodity (ATA) disks for very short-term storage. This file system currently sustains around 250MB/s.

#### **C.10.3.1.2 Washington, DC**

The home file systems on NOAA's HPCS system located at Washington DC consist of 25.6 TB of disk space. Most of this disk space is configured as a General Parallel File System (GPFS), which is integrated with the switch as described in section C.10.1.2 above. GPFS utilizes 2 GB Fibre Channel adapters per storage and I/O LPAR.

#### **C.10.3.1.3 Princeton, NJ**

In Princeton, NOAA uses a unified /home filesystem for all scientific computers and workstations. This allows for users to log into various computational resources and have their environment and files follow them. This has required custom .cshrc and .login dotfiles that are designed for the unified home directory.

By default, each scientific user is limited to 10 GB of home directory disk space. /home is used primarily for model source code, batch job printouts, and files for workstation applications. Large data files are kept in /archive (the filesystem that is under HSMS).

The home file system is a 2.3TB SAN from a SGI Origin 3000 server. This file system is being served via CXFS within the HPCS and via NFS to the scientific workstations. Currently, the filesystem is 60% used by 10 million files. Backups are currently being done via xfsdump to the STK 9940B tapes in the silos.

### **C.10.3.2 Hierarchical Storage Management System (HSMS)**

#### **C.10.3.2.1 Boulder, CO**

NOAA's HSMS in Boulder is based upon ADIC's StoreNext software. The robotic component is an ADIC AML/J with 8 LTO tape drives. Managed file systems (disk cache) use the DDN S2A 6000s mentioned above. Currently, 1 TB of space is allocated for two managed file systems. Over 88 TB is managed and is composed of over 4,000,000 files. Access to the HSMS is primarily via locally written get/put scripts, some limited direct access to the managed file systems is allowed.

#### **C.10.3.2.2 Washington, DC**

NOAA's HPCS in Washington DC has access to a 2 PB HSMS. The HSMS is presently considered part of the operational system at NCEP. This storage

system, in principle, consists of tape storage with dedicated disk storage for staging data to be stored to, or retrieved from tape. This system is based on HPSS. Access to this system is through htar and PFTP commands, and a public domain HSI Unix-like interface. Additional information on HPSS and HSI can be found at:

<http://www.hpss-collaboration.org>

<http://www.sdsc.edu/Storage/hsi/>

### **C.10.3.2.3 Princeton, NJ**

NOAA's HSMS in Princeton uses SGI's Data Migration software to provide a data archive via the /archive filesystem. This software manages both a disk and a tape copy of each file in /archive. This is done transparently to the user, under one file name.

The HSMS software has a 15.8 TB SAN for file staging. The disk SAN is connected to the rest of the HPCS by sixteen 2 Gb FC per node. There is currently 2.5dPB stored in 10.7 million files on 9940B (200GB/tape) and 9840 (20dGB/tape) tapes. There are also 2.4 million disk-resident files that are smaller than 64 KB.

To archive a file, a user simply copies or moves the file into their /archive/<user> directory from any node in the HPCS. Once a new file is present in /archive, data migration will automatically make a tape copy of the file. At first, the disk copy of the file is also kept available. If the file is not accessed for several days, data migration will remove the disk copy of the file, keeping only the tape copy. The next time this file is accessed, it will be "staged" from tape to disk. Subsequent accesses will then be from the disk-resident copy of the file. However, files may be moved between subdirectories, renamed, or removed without causing the file to be "staged" from tape to disk.

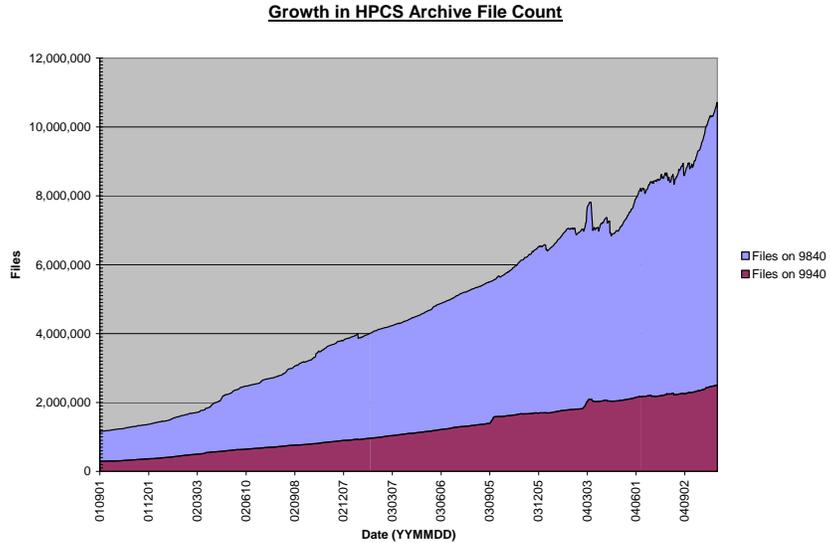


Figure 2 - Growth in tape-resident files on HSMS in Princeton

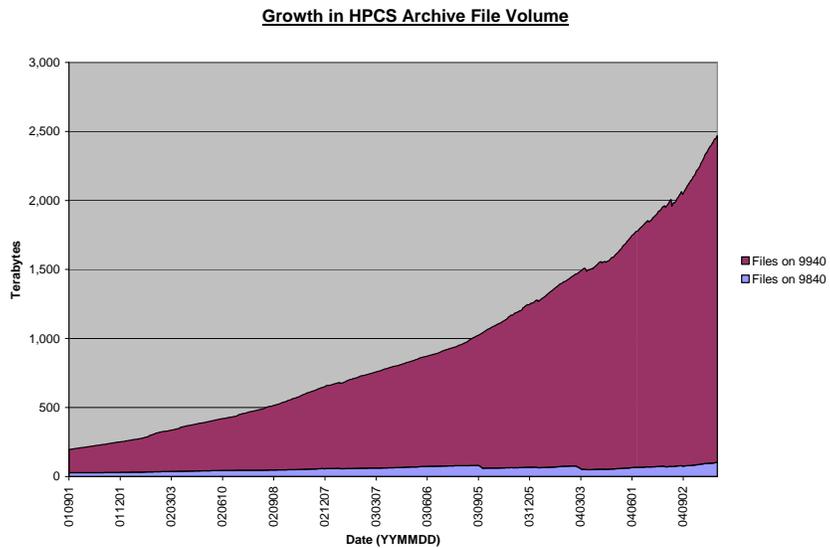


Figure 3 - Growth in tape resident storage on HSMS in Princeton

### C.10.4 Networking

#### C.10.4.1 Wide Area Networks (WANs)

	Internet	Internet2	Other
Boulder	155 Mb/s	310 Mb/s	NLR is expected Q3FY2005
Washington, DC	622 Mb/s	155 Mb/s	
Princeton	9 Mb/s	100 Mb/s	1Gb/s – Princeton University Sayre Hall (Forrestal campus)

**C.10.4.1.1 Boulder, CO**

NOAA's Boulder facility is connected to the Abilene network via the Front Range GigaPOP at 310Mb/s. A 155Mb/s connection is available to Boulder through a commercial Internet provider. By the third quarter of Fiscal Year 2005, Boulder will be part of the National Lambda Rail (NLR) with connectivity of at least 1000Mb/s.

**C.10.4.1.2 Washington, DC**

NOAA's Camp Spring, MD facility is connected to Internet2 with an OC3. An OC12 internet link is presently used close to capacity. There is some capacity left on the OC3 link. This setup is propagated from the previous operational use of the present HPCS in Washington, DC. It is expected that future HPCS systems will rely more on Internet2.

**C.10.4.1.3 Princeton, NJ**

NOAA's Princeton facility is connected to Princeton University's Sayre Hall Building via 1Gb/s Ethernet. NOAA connects to Princeton University through a 100Mb/s microwave link that is accessed through the Sayre Hall Building. The 100 Mb/s Internet2 connection traverses the Princeton University connection. NOAA connects to the Internet through a 9Mb/s connection provided by a commercial Internet provider.

A Cisco 2600 router and a Juniper M7i router are used to connect to the commodity Internet, Princeton University, and subsequently Internet2. These exterior border routers, along with the local DMZ networks, are segmented off from the LAN by redundant Cisco PIX 535 firewalls.

**C.10.4.2 Local Area Networks (LANs)****C.10.4.2.1 Boulder, CO**

NOAA's Boulder LAN backbone that connects to the HPCS is based upon Gigabit Ethernet.

**C.10.4.2.2 Washington, DC**

NOAA's Washington, DC LAN that connects to the HPCS consists of Cisco Catalyst 3500 XL routers at 100 Mb/s. Apart from several hundred desktop systems, the LAN includes several servers. Servers are typically high-end SGI Origin 300 or SGI Altix servers which are used for storage, post-processing and development.

**C.10.4.2.3 Princeton, NJ**

NOAA's Princeton LAN backbone that connects to the HPCS consists of Brocade, Cisco, Enterasys and Juniper equipment. The HPC backbone ethernet network is handled by a Cisco Catalyst 6509. The HPC nodes are

connected at both 100 Mb/s and 1Gb/s. To support the large data requirements of the system, Brocade SilkWorm and Cisco Catalyst switches are used to connect the SAN infrastructure to the HPC nodes. Local users connect to the LAN via distribution switches at either 100Mb/s or 1Gb/s.

NOAA's Storage Area Network (SAN) in Princeton has 22.5TB of capacity that is accessible through switched 2 Gb/s FC. The LSC in Princeton has eight 2 Gb/s channels per node to this SAN. The AC in Princeton has twenty-four 2 Gb/s channels per node to this SAN.

Basic network services provided include: DNS, LDAP, Mail, NFS, NIS, NTP, and Printing.

### **C.10.4.3 Batch queuing systems**

#### **C.10.4.3.1 Boulder, CO**

NOAA's HPCS in Boulder uses the open source version of Sun Grid Engine (SGE). Currently version 5 is in use on the primary set of clusters. Version 6 will be in use by November 2004. With version 5, a custom prescheduler is used to manage jobs based upon the account specified by the user. Accounts are given a maximum number of CPUs to be in use at any given point in time, a maximum number of jobs in use at any given point in time, a maximum job length and maximum priority. Certain accounts are also given dedicated resources (this will be replaced by resource reservations with version 6). Accounts are members of classes (classes can be members of classes as well), classes have total resource limits as well that the prescheduler takes into account. The qsub command is "wrapped" with a custom script that verifies that the user is a member of the account and hasn't exceeded any of the per job maximums. The "wrapper" also modifies priority 0 jobs to use an alternate backfill account associated with the primary account. Typically at least 70,000 jobs per week pass through the batch system.

#### **C.10.4.3.2 Washington, DC**

NOAA's HPCS in Washington, DC uses LoadLeveler. The system is set up with separate partitions for serial and parallel jobs (see section C.10.2.2 above). Queues have been assigned to provide access to specific job resource requirements and priority. Accounting is presently performed for system monitoring only, but is deemed essential in future R&D resource allocation.

#### **C.10.4.3.3 Princeton, NJ**

NOAA's HPCS in Princeton uses Grid Engine Enterprise Edition 5.3. This is a POSIX 1003.2d batch environment (with extensions) that has been provided by Raytheon. IRIX cpuset support has been added. More information on the batch queuing system at <http://gridengine.sunsource.net/>

Grid Engine projects are used to track usage and assign higher job scheduling priority to some activities. Only users designated by their team may use these projects.

Each calendar month, a running total is kept of each group's CPU usage. Each group is assigned a monthly CPU allocation by the GFDL Director. When a group has exceeded its monthly allocation, that group's jobs are blocked from running, however the group is still allowed to run jobs in the windfall queue.

## **C.10.5 Software**

### **C.10.5.1 COTS**

#### **C.10.5.1.1 Boulder, CO**

NOAA's HPCS in Boulder currently has licenses to use IDL on the visualization servers. The HPCS in Boulder also has licenses for the Portland Group Fortran and C compilers as well as the Intel Fortran and C compilers.

#### **C.10.5.1.2 Washington, DC**

NOAA's HPCS system located at Washington DC presently has licenses to use the following COTS software on the LSC: TotalView, Vampir, IDL, Matlab. Also available are the following libraries: Visual Numerics IMSL, ScaLaPack.

#### **C.10.5.1.3 Princeton, NJ**

NOAA's HPCS in Princeton currently has licenses to use the following COTS software on the HPCS: CASEVision/WorkShop, MIPSpro f90 Compiler, Etnus TotalView, F-lint, IDL, Legato Networker, Maple, Mathematica, Matlab, MIPSpro f90, NAG Fortran Libraries, NAG Iris Explorer, and S-Plus.

### **C.10.5.2 Community Supported Software**

#### **C.10.5.2.1 Boulder, CO**

NOAA's HPCS in Boulder currently utilizes the following open source software: SUSE Linux, Fedora Linux, Ferret, NCAR graphics, GrADS, and NetCDF utilities.

#### **C.10.5.2.2 Washington, DC**

NOAA's HPCS system located at Washington DC presently utilizes the following open source software: GrADS, NCAR graphics, ImageMagick, NetCDF utilities, GEMPAK, and ISMF.

#### **C.10.5.2.3 Princeton, NJ**

NOAA's HPCS in Princeton currently utilizes the following open source software: Ferret, Grace, GrADS, NCAR graphics, and NetCDF utilities.

## **C.10.6 Reliability, availability and support**

### **C.10.6.1 Boulder, CO**

The current contractual uptime requirement for the HPCS in Boulder is significantly exceeded by the incumbent. For carefully monitored jobs, the reliability of the system in terms of delivery of results to an end user is over 98.5%.

The current facility has 250kVA available to the current HPCS and commensurate cooling.

The incumbent supplies two on-site software engineers to support the facility. A program manager is on-site over 50% of the time. Other software engineers work remotely on the system a small (<10%) percentage of time to provide additional support.

### **C.10.6.2 Washington, DC**

#### Reliability:

Reliability is expressed in terms of Mean-Time-Between-Failure (MTBF). A failure is defined as any scheduled or unscheduled event that triggers down time. Individual components of the system can fail and not be counted as a system failure if the system has sufficient redundancy, error correction or resiliency to allow the failure to not affect individual jobs running in production or development.

#### Availability:

Availability is determined by computing the ratio of total batch computation processor hours available for execution of the operational and developmental suites to the total batch computation processor hours each month. The time a batch computation processor is available for execution is determined by subtracting processor down time from wall-clock time. Spare nodes or processors may be configured into the batch computation processor pool upon failure of a node and/or processor in the batch computation pool to reduce down time, but down time accumulates until the spare is made available in the system for execution. Failure(s) of other components of the HPCS that result in batch computation processor(s) being unable to run jobs is equivalent to failure of those batch computation processors themselves.

#### Down Time:

Down time is that period of time when the HPCS or a component thereof is inoperative. Down time shall commence at the time when NOAA contacts the Vendor's maintenance representative at the designated point of contact or the Vendor's answering service or other continuous telephone coverage provided to

permit NOAA to make such contact to report a failure. Down time shall end when the HPSCS is returned to NOAA in operable condition. The HPSCS, or individual component thereof, may be declared inoperative while problem diagnosis takes place. Down time includes any time required for operating software regeneration, reinstallation, or reconfiguration. During a period of down time, NOAA may continue to use operable components of the HPSCS when such action does not interfere with maintenance of the inoperative components of the HPSCS.

Support:

The Government requires 7x24x365 System Administration support. Current staffing levels are at 3 System Administrators, 1.25 Application analysts and one part time HSM specialist.

**C.10.6.3 Princeton, NJ**

The HPC System continues the Government's historically high utilization of its computing resources. Reliability, availability, and Contractor support are considered fundamental aspects of the HPSCS

Downtime

Downtime is that period of time when all of an HPSCS component's workload cannot be accomplished due to a malfunction in the Contractor-maintained equipment or software, or when the HPSCS or a component of the HPSCS is released to the Contractor for maintenance. Periods of Remedial and Preventative Maintenance count as downtime.

Null time is that period of time when the workload cannot be accomplished due to circumstances beyond the scope of the contract. Null time does not count as downtime.

Downtime for each HPSCS component is based on the fraction of the resources available for that component's workload, arrived at through consultation between the Government and the Contractor, and determined ultimately by the Government. Downtime is accumulated on the LSC, Post-Processing and Analysis computers if they are not able to perform the workload when either the HSMS or HFS is down.

The Contractor provides the Government a designated point of contact to request maintenance. The Contractor maintains escalation procedures that allow the Government round-the-clock telephone contact with knowledgeable Contractor staff should the designated point of contact be unavailable.

A component's downtime shall commence at the time the Government makes a bona fide attempt to contact the Contractor at the designated point of contact. At this time the Government begins a log of the problem. Information to be entered into the log is determined by the Government.

A component's downtime excludes any time in which the Government denies the Contractor maintenance personnel access to the malfunctioning hardware and/or software.

A component's downtime ends when the computer is returned to the Government in operable condition as determined by the Government, ready to perform all of the workload.

Preventative Maintenance is completed outside of primetime hours (7 am – 7 pm).

The testing and installation of every major operating system release installed at the request of the Government and one (1) minor operating system release installed at the request of the Government during any annual period will count as downtime.

Preparation for the execution of post-upgrade LTDs, including any benchmark development by the Contractor, associated with the agreed-upon upgrades proposed at contract award, will count as downtime.

#### Availability

Proposed throughput benchmark performance levels are combined with the proposed availability level to determine a measure of overall proposed system-life throughput for the LSC and for the Analysis Computer. The actual throughput will be measured on a periodic basis, to be determined by the Government and Contractor, by combining the demonstrated benchmark performance with the operational use time on the LSC and on the Analysis Computer. The proposed performance levels must be met for each measurement of actual throughput regardless of past delivery of suites.

Shortfalls in throughput on the LSC or on the Analysis Computer are made up with new equipment brought in at no additional cost to the Government. Using the demonstrated benchmark performance on the upgraded LSC or Analysis Computer, the Government calculates how long the upgrade shall stay in place to compensate for the shortfall in throughput. This calculation is rounded up to a multiple of 6-month intervals to minimize disruption.

At the option of the Government, shortfalls in throughput on the LSC or on the Analysis Computer due to downtime shall cause downtime credits to accrue. These downtime credits shall be in lieu of bringing in new equipment. Downtime credits shall accrue on the HSMS or HFS.

To better reflect the Government's computational needs over time, changes in the LSC and AC throughput benchmarks are made by mutual agreement between the Government and the Contractor throughout the life of the HPCS.

Accumulated computational cycles (in CPU-hours) that are lost when jobs are lost due to component failure or component reboot are not counted toward the system-life throughput calculation. If the accounting software cannot report the

accumulated computational cycles, it will be assumed that 4 CPU-hours were lost for each processor on which a job ran.

All performance levels proposed for hardware and software upgrades must be met regardless of past delivery of suites

Uninterruptible power supplies (UPS) is required for all components of the HPCS. The UPSes provide sufficient power during environmental failures to gracefully shut down all components of the HPCS.

#### Support

The Government requires at least two software engineers (to provide a comprehensive system administration service), at least one hardware engineer, and at least one applications analyst available on site for at least eight hours within GFDL's primetime window, five days per week. Additional on-call support is provided 24 hours per day, 7 days per week, with a 2-hour response time. The Government reserves the right to substitute hardware engineers with software engineers or application analysts during the life of the contract on an as-needed basis.

The Government requires an itemized list of all Contractor-supplied hardware and software items, and documentation of these items.

### **C.11 Appendix B - Available Government Facilities**

This document provides descriptions of three government-furnished facilities that are designed to support high-performance computer systems and are available for possible use by Offerors to house systems under the NOAA HPC R&D contract, subject to the availability of facility resources as projected in Section C.8.11 below under the assumptions provided therein and subject to the facility terms and conditions provided in Section C.5.1. Two of these facilities, designated as BLDR-1 and BLDR-2, are located in the David Skaggs Research Center (DSRC), 325 Broadway, Boulder, CO, and the third, designated as PRTN, is located in Geophysical Fluid Dynamics Laboratory (GFDL) building complex, referred to as the "Princeton Complex" below, at 201 Forrestal Road, Princeton University Forrestal Campus, Princeton, NJ.

BLDR-1 currently houses NOAA's JET computer system operated by OAR's Forecast Systems Laboratory under a contract with HPTi. The BLDR-1 facility will not be available until October 2006. BLDR-2 is a facility currently under construction at DSRC and is expected to be available for use in October 2005. PRTN currently houses NOAA's HPCS operated by OAR's Geophysical Fluid Dynamics Laboratory under a contract with Raytheon.

#### **C.11.1 Layout and Physical Dimensions of Computer Room**

##### **C.11.1.1 Computer Room Layout**

##### **BLDR-1**

Figure 1 shows the computer room layout for BLDR-1. The area within the dotted line will be available in October 2006. This room is only designed for medium-density cooling configurations.

#### **BLDR-2**

Figure 2 shows the planned computer room layout, which is currently in the design phase. It will be designed with both overhead and under floor extreme-density cooling. The uninterruptible power supplies (UPS) will be fed from utility power only (i.e., the room will not be connected to the emergency motor generator). This room will be completed and available for use by the Offeror by October 2005.

#### **BLDR-1 & BLDR-2 Building Diagrams**

Offerors requesting building diagrams or related documents of federal government facilities shall be required to have an authorized representative execute a non-disclosure form as required under GSA Order PBS 3490.1. Specific room floor plans are available upon request and are not subject to these regulations.

#### **PRTN**

Figure 3 shows the computer room layout for the current system, but with the following planned modifications: move silos onto the northern hardpan and construct five office cubicles on the southern hardpan. References to the front of the room in the following discussion refer to the bottom of the figure (nearest to the Operators Room), while the back of the room is at the top of the figure.

Figure 4 (attached) shows the overall Computer Building layout. The rooms shown at the bottom of the figure from left to right are:

- Loading Dock, which is designed to accept deliveries from 18-wheel trucks.
- Storage Room adjacent to the Loading Dock, which also serves as a receiving/staging area for deliveries to the Laboratory.
- Vendor Room, which currently provides office space for Contractor personnel of the existing system (This may be converted into office space with vendor offices relocated to the southern hardpan.)
- Operators Room, which serves as the control room for computer operations as well as security monitors for the Computer Room as well as outside building access points for the Computer Building and Main Building.
- Printer/User Output Room, which contains local Computer Building printers as well as user output bins.
- Operations Lounge.

#### **C.11.1.2 Location of Computer Room and Characteristics of Surrounding Campus**

##### **BLDR-1 & BLDR-2**

Both BLDR-1 and BLDR-2 are located within the David Skaggs Research Center (DSRC) in Boulder, CO. Due to the large number of visitors, the following website

was developed to assist with locality information with regard to the site: <http://boulder.noaa.gov/> This website will display maps, local information, driving directions, etc. Regarding specific locations: BLDR-1 is located in the “B” Block of the DSRC, on the 2nd floor, and BLDR-2 is located in the “A” Block of the DSRC, in the Garden level (basement).

#### **PRTN**

The Computer Room is in the Computer Building, which is one of two buildings in the PRTN Complex located at 201 Forrestal Road, Forrestal Campus, Princeton, NJ 08540 on Site B of Princeton University’s Forrestal Campus. This campus is currently devoted to university research and is intended by Princeton University to be developed in the future as an office park on the U.S. Route 1 corridor. The nearest airport is the Trenton-Mercer Airport, which is roughly 12 miles away. The nearest highway, U.S. Route 1, is roughly 1/3 mile away.

### **C.11.1.3 Physical Dimensions of Computer Room**

#### **BLDR-1**

Figure 1 indicates the computer room layout for BLDR-1. The entire computer room is 3600 square feet in size, with dimensions of 60 feet by 60 feet. The area within the dotted line in the figure will be available in October 2006 and is estimated to be approximately 2250 sq. ft.

#### **BLDR-2**

Figure 2 indicates the computer room dimensions and layout for BLDR-2. The entire computer room space is approximately 2100 sq. ft. A portion of this space is used for a substantial, Americans-with-Disabilities-Act (ADA)-compliant ramp for room access. 1500 sq. ft. of raised floor space is available for equipment. An adjacent room contains an additional 325 sq. ft. of raised floor space, separated from the main room by a wall.

#### **PRTN**

Figure 3 indicates the computer room layout for PRTN. Figure 4 shows the location of the computer room within the Computer Building of the PRTN Complex. The entire computer room is 10,004 square feet in size, with dimensions of 122 feet by 82 feet; this includes the UPS Room, which is located on the hardpan on the right rear corner of the Computer Room.

### **C.11.1.4 Raised Floor Space**

#### **BLDR-1**

A total of 3600 sq. ft. of raised floor is in BLDR-1. 2250 sq. ft. of this space will be available in October 2006.

#### **BLDR-2**

A total of 1500 sq. ft. of raised floor will be in BLDR-2, all of which will be available once construction of the room is completed in October 2005.

#### **PRTN**

The raised floor area in PRTN totals 7052 sq. ft. with dimensions of 86 feet by 82 feet.

**C.11.1.5 Non-Raised Floor Space and Equipment Staging Areas****BLDR-1 & BLDR-2**

There is no non-raised floor in either computer room. The computer rooms themselves are used as staging areas.

**PRTN**

The non-raised floor on the left side of the room is 1476 sq. ft. with dimensions of 18 by 82 feet. The non-raised floor on the right side is 936 square feet, with dimensions of 18 by 52 feet, reflecting reduced non-raised floor space due to the presence of the UPS Room. Figure 3 indicates the planned future location of the five (5) StorageTek silos, four of which are currently located on the raised floor but will be relocated this winter to the locations on the hardpan as indicated in the figure.

**C.11.1.6 Space for Vendor Personnel, Maintenance Space, and Vendor Storage****BLDR-1 & BLDR-2**

The Government will supply two (2) 150-sq. ft. offices and approximately 100 square feet of lab space for use by the Contractor. None of this space will have furniture or lab benches. The space will have up to four (4) telephones as well as adequate LAN connections.

**PRTN**

Vendor support personnel (2 Computer Engineers and 2 Software Engineers) currently occupy three (3) rooms in the Computer Building. Equipment storage and maintenance space is provided on the left hardpan in the Computer Room. Plans are being made to construct five (5) offices on the left hardpan in the computer room to provide additional vendor office space while freeing up one of the current vendor rooms for other use.

**C.11.1.7 *Height of ceiling above raised floor*****BLDR-1**

Drop ceiling height = 8.5 feet

**BLDR-2**

Drop ceiling height = 10 feet

**PRTN**

Drop ceiling height = 9.5 feet

**C.11.1.8 *Maximum Allowable Height of Equipment***

The following equipment height are mandated by fire codes and cooling limitations.

**BLDR-1 & BLDR-2**

Due to fire code and cooling limitations, the maximum allowable rack height will be 42U.

**PRTN**

Due to fire code regulations, the maximum allowable rack height will be no more than 18” below the drop ceiling height.

#### **C.11.1.9      *Space Below Raised Floor***

The space below the raised floor is defined as height from the sub-floor base to the top of the raised floor.

##### **BLDR-1**

Raised floor height = 12 inches

##### **BLDR-2**

Raised floor height = 24 inches

##### **PRTN**

Raised floor height = 24 inches (except for the deeper chilled water trench, 4 feet deep, that extends down the middle of the room as indicated in Figure 4).

### **C.11.2      Installation Characteristics and Raised Floor Capacity**

#### **C.11.2.1      *Physical Access for Equipment Installation***

##### **BLDR-1 & BLDR-2**

Loading Dock - Standard, designed to accommodate 18-wheeled semi-trucks.

Access Path –

Corridors - 8” raised floors and 7-ft high doorframes

Elevator – Freight Elevator with 8000-lb. Capacity

BLDR-1: 25’ from Loading Dock to the freight elevator and then 50’ to the computer room

BLDR-2: 25’ from the Loading Dock to the freight elevator and then 150’ to the computer room, down and up ADA-compliant ramps.

##### **PRTN**

Loading Dock – Standard, designed to accommodate 18-wheeler semi-trucks.

Access Path – Equipment passes through an entry door from the Loading Dock into the Storage Room and then onto a hardpan staging area in the south corner of the Computer Room - a total distance of roughly 30 feet. The two sets of double doors along this path have clearances of 85 inches high by 70 inches wide.

#### **C.11.2.2      *Staging and Assembly Areas***

##### **BLDR-1 & BLDR-2**

All staging and equipment assembly must be performed within the computer room.

##### **PRTN**

Equipment may be staged and assembled on the hardpan in the computer room immediately adjacent to the entry from the Storage Room and the Loading Dock. This non-raised floor area is at least 18’ wide by 20’ deep.

**C.11.2.3 Loading Capacity of Raised Floor****BLDR-1**

Computer Room Raised Floor Specs: ConCore SF 1250 Bolted Stringer:  
Concentrated Load: 1250 lbs.; Uniform Load: 300 lbs./ft<sup>2</sup>; Ultimate Load: 3850 lbs.; Rolling Load: 1000 lbs. (10 Passes)

Raised floor Specs (adjacent hallways): ConCore SF 1250 Cornerlock:  
Concentrated Load: 1250 lbs.; Uniform Load: 300 lbs./ft<sup>2</sup>; Ultimate Load: 3750 lbs.; Rolling Load: 1000 lbs. (10 Passes)

Floor tiles are 2' x 2' and are shown by location in the Figure 1 drawing. The dotted squares represent perforated tiles. Network and power connections are made via 5" round ports within the solid tiles only. All tiles, unencumbered by racks or other equipment, can be relocated for improved airflow, if necessary.

**BLDR-2**

Computer Room Raised Floor Specs: ConCore SF 1250 Bolted Stringer:  
Concentrated Load: 1250 lbs.; Uniform Load: 300 lbs./ft<sup>2</sup>; Ultimate Load: 3850 lbs.; Rolling Load: 1000 lbs. (10 Passes)

Raised floor Specs (adjacent hallways): Concrete floors

**PRTN**

The raised floor in the Computer Room is designed to support a uniform live load of 250 pounds per square foot, with a deflection of not more than 0.040 inch. The raised floor is electrically interconnected to provide a common electric reference. Most of the raised floor, indicated as 2'-by-2' square tiles in Figure 3, was replaced as part of the site preparation for the current Raytheon contract. The exceptions for this raised floor upgrade were for tiles under the following equipment: the two oldest StorageTek silos (when located on the raised floor), the MGE 250-kVA UPS, the Dataflow air-handler along the rear wall of the Computer Room, the Liebert 125-kVA PDU, and the EPE 125-kVA PDU.

Adjacent Non-Raised Floor and Pathway to Loading Dock: Concrete floors

**C.11.3 Power Facilities****C.11.3.1 Power Service to Building****BLDR-1 & BLDR-2**

480/277-volt main switchboard, rated for 3000 amps and 65,000 AIC. The gear has integrated IMPACC monitoring software. All power is fed from a single utility substation. Power is then distributed via step-down transformers and panel boards.

**PRTN**

The Princeton Complex receives its power from Public Service Electric and Gas (PSE&G) Company. A PSE&G utility substation with a 2500-kVA, 13200-to-

4160-volt, oil-filled transformer and 4160-volt switchgear provides electrical service to the Princeton Complex. This substation is located outside of the southwest corner of the Computer Building directly adjacent to the parking lot. This substation provides power support to the Princeton Complex as well as to several of the other buildings on the B site of the Forrestal Campus. Power usage for these other buildings is primarily for offices, although research activities in some of these other buildings occasionally require substantial power from the substation. The power load shown on the meter servicing these other buildings recently showed a load of 1070 kVA. PSE&G has indicated that it will upgrade the capacity of this substation, if required; however, the underground feeder line connecting the substation to the main PSE&G grid (located on nearby U.S. Route 1) would be an additional cost to the customer.

An underground 4160-volt feeder, dedicated to the Princeton Complex, is routed from the utility substation to separate building substations located within the Princeton Main Building and Computer Building. The Main Building substation (1500-kVA 4160-to-480-volt transformer), located on the ground floor of Princeton's Main Building, provides power to the Main Building and to the Chilled Water Plant for the Princeton Complex. A recent reading from the Main Building substation showed a peak load of 1000 kVA, which probably occurred as a result of peak cooling load during the summer.

The Computer Building substation is located in the Transformer Room, the location of which is shown in Figure 4. This substation is comprised of a 4160-volt air interrupter switch, a 1500-kVA silicone-filled transformer, and a 2000-ampere, 480/277-volt main switchboard. The equipment was installed around 1980 when the Computer Building was constructed and provides power to the Computer Building. A recent reading from the Computer Building substation showed a peak load of 540 kVA.

The lighting, large mechanical equipment (pumps, A/C units, etc), and some computer equipment are served at 480/277 volts via panel boards located throughout the Computer Building. The building receptacles, small motors, desktop computers, computer room equipment, and similar loads are served at 208/120 volts via step-down transformers and panel boards. Currently all three chillers, described below, use the Main Building substation for power, a configuration that restricts their joint operation. Installation of an A-B switch for one or more of these chillers to permit use of power from the Computer Building substation, is a possible option for increasing available cooling capacity and providing increased flexibility for N+1 chiller backup.

### ***C.11.3.2 Cost of Electrical Utilities Based on Recent Billing***

#### **BLDR-1 & BLDR-2**

Electrical utility costs for DSRC occupants are currently charged to building tenants according to finished square footage occupied. Based on this algorithm, the electrical cost was \$3.14 per square foot in FY 2004. There is a real

possibility of metering electrical usage in the future. In this case, electrical costs would increase substantially. In FY 2004, the average cost for power was \$0.05/kWH.

### **PRTN**

The following table shows usage and expenditure data for the entire Princeton Complex as an indicator of recent electrical utility costs.

Total Usage and Expenditures for Electrical Utilities of the Princeton Complex  
for FY 1999 - FY 2004

	Annual Usage (kWH)	Expenditure	Avg. Cost / kWH
FY 1999	8,826,400	\$713,948	\$0.0810
FY 2000	9,032,000	\$640,544	\$0.0707
FY 2001	8,555,200	\$644,035	\$0.0756
FY 2002	5,976,000	\$431,506	\$0.0722
FY 2003	6,593,600	\$472,405	\$0.0711
FY 2004	7,981,571	\$654,619	\$0.0820

Measurements taken from the two UPS systems in the Computer Room suggest that the power usage of the current HPCS is roughly 400 kVA. If we assume that the power required for cooling to support this system is conservatively estimated to be 1.2 times this power usage (accounting for other room heat sources and cooling system inefficiencies), this implies a total power usage for the system of around 880 kVA. Assuming this power usage is constant throughout the year implies an annual usage of roughly 6,500,000 kWH or roughly 82% of the total usage for the Princeton Complex.

### **C.11.3.3 Power conditioning and UPS capabilities**

#### **BLDR-1**

- Power Distribution: Cutler-Hammer Electrical Distribution Equipment (480Volt, 3 Phase) Under floor distribution is accomplished by 50 ft. flexible conduit power whips, fed from wall-mounted breaker panels.
- UPS: 300 kVA Chloride UPS Systems (450 kVA Installed); (250kVA available to R&D HPCS in 2006); 8-Minute Runtime (Full Load)
- Other: Transient Voltage Surge Suppressor (TVSS) Protected
- Emergency Power Off (EPO) Switch Protected

#### **BLDR-2**

- Power Distribution: Cutler-Hammer Electrical Distribution Equipment (480Volt, 3 Phase)
- Under floor power distribution will be accomplished by fixed receptacles attached to rigid conduit and mounted one foot below the floor tiles.

- UPS: 350 kVA UPS System (500 kVA Installed); 16-Minute Runtime (Full Load)
- Other: Transient Voltage Surge Suppressor (TVSS) Protected
- Emergency Power Off (EPO) Switch Protected

**PRTN**

- Power Distribution Units (PDUs) [shown in dark blue in Fig. 3]:
- Three (3) United Power 225-KVA PDUs (owned by Raytheon)
- Liebert 125-KVA PDU (owned by Government/Princeton University)
- EPE 125-KVA PDU (owned by Government/Princeton University)
- UPS [shown in brown in Figure 3]:
- 500-kVA UPS MGE cabinet and battery bank (owned by Raytheon)
- 225-kVA UPS MGE cabinet and battery bank (owned by Government/Princeton University)

**C.11.3.4 Backup Power Generator Capabilities****BLDR-1**

- All Equipment fully backed up by emergency backup generators.
- Current Generator is 1250-KW Cummins Diesel. 1800-gallon tank capacity allows 24 hours of operation

**BLDR-2**

The DSRC EM Generator has reached its rated capacity, and no further load will be allowed. Therefore, the BLDR-2 facility will not have EM generator backup. The UPS systems will be designed to sustain a 16-minute outage at full load. The recent power loss history at DSRC is presented in Section 8a below.

**PRTN**

Natural-gas-fired, 75-kW generator: activates automatically when power is lost. This generator is only designed to allow pumps within the chilled-water cooling system to continue to run during a brief power outage.

**C.11.4 Cooling Facilities****C.11.4.1 Cooling Service to Building****BLDR-1 & BLDR-2**

Building operates three (3) 470-ton chilled water cooling systems. Two of the systems are on EM generator power, and one is fed strictly from utility. There are three cooling towers, each sized to match the capacity of a single chiller. There are three primary pumps, three secondary pumps, and three condenser water pumps. The chiller plant is located in a mechanical room that is adjacent to BLDR-2. Cooling is delivered at 42°F. There is a rough-in plan for an additional chiller and space for an additional cooling tower in the future, but no plans have been authorized to expand the chiller plant at this time. The chiller plant is currently operating at full capacity.

**PRTN**

Primary Chilled Water Plant, located in the mechanical room and tower bay southeast of the Transformer Room (see Figure 4), contains two chillers:

- 400-ton Carrier centrifugal chiller (designated Chiller #1)
- 350-ton York centrifugal chiller (designated Chiller #3)
- Baltimore Air Coil Cooling Towers

Chiller #1 was installed in the spring of 2000 along with new cooling towers and pumps. Chiller #3 was installed in 1996.

Secondary Chilled Water Plant, located in the Transformer Room (see Figure 4):

- 225-ton Carrier centrifugal chiller (designated Chiller #4)
- Baltimore Air Coil Cooling Tower

Chiller #4 was installed in 1979. Cooling tower replacement will be completed by November 2004.

The Primary Chilled Water Plant provides cooling to the entire Princeton Complex. The Secondary Chilled Water Plant is designed to only support the Computer Building.

The current configuration only permits operates with only one chiller in the Primary Chilled Water Plant to operateoperating at one a time. The power capacity of the Main Building substation has been determined to permit operation of any two of the three chillers at one time. The issue of whether the two large chillers, Chiller #1 and #3, can be run simultaneously requires further study; however, indications are that they cannot, due to water flow constraints caused by the current chilled water piping designThese restrictions are due to limitations on power usage (i.e., all chillers run off of the Main Building substation – see above) and to flow restrictions caused by the current chilled piping design. If water flow constraints are the problem, as is expected, then these restrictions could be eliminated by a redesign of the chilled water plumbing; until this situation is resolved, operation will be limited to either Chillers #1 and #4 or Chillers #3 and #4 for heating loads requiring more than one chiller. The problem with this is that, in the event that Chiller #4 is unavailable, operation will be limited to only one chiller, either Chiller #1 or #3..

Because of problems with the cooling tower in the Secondary Chilled Water Plant, Chiller #4 has only been used on especially hot days for emergencies, or as backup to partially support the cooling load during periods of time when one of the primary chillers is taken off-line for servicing or repair. However, the cooling tower for the Secondary Chilled Water Plant has been replaced during the fall of 2004. Because of this enhancement, Chiller #4 is now expected to be capable of providing sustaining a more robust cooling load.

Cooling is delivered to the Computer Room through a six-inch and eight-inch piping system from the mechanical room at a temperature of 42° F, plus or minus 2 degrees. The pipe enters the computer room in a trench that is 4 feet deep under the raised floor in the center of the computer room, as indicated in

Figure 4. The piping to the Computer Building currently limits the flow of chilled water to the Computer Room and will require redesign if additional cooling capacity is required.

#### **C.11.4.2 Room Air Conditioning Capabilities**

##### **BLDR-1**

- 90-Ton Liebert Downdraft Computer Room Air Conditioners (CRACs) (De-rated for altitude) (120 Tons Installed via 4 x 30 Ton units)
- Cooling powered from Emergency circuits.

##### **BLDR-2**

- 90-Ton Liebert Downdraft CRAC (De-rated for altitude) (120 Tons Installed via 4 x 30 Ton units)
- Thirteen (13) 4-Ton Liebert Extreme Density Overhead (XDO) Systems
- Cooling powered from UPS

##### **PRTN**

- Seven (7) 35-ton Dataflow/APC CRACs (owned by the Government/ Princeton University)
- Other smaller air handling units - located in the ceilings of the Printer Room and the UPS room

The primary source of humidification for the Computer Room is an air conditioner and steam boiler, located in the building tower on the western corner of the Computer Building.

#### **C.11.4.3 Availability of plumbing for chilled water**

##### **BLDR-1 & BLDR-2**

All chilled water plumbing is installed to accommodate the Computer Room Air Conditioner (CRAC) units and the Extreme Density Overhead systems. No additional plumbing capability or capacity is available.

##### **PRTN**

The computer room has operated several water-cooled computer systems (most recently - Cray T932, T94, and T3E) and thus has been previously configured to support water-cooling, including a chilled-water trench running down the middle of the raised-floor area of the computer room, as indicated in Figure 4.

### **C.11.5 Networking Facilities**

#### **C.11.5.1 Wide-Area Networking Services To Computer Room**

##### **BLDR-1 & BLDR-2**

NOAA Boulder maintains multiple wide-area networking links. These include a one-gigabit per second (1 Gb/sec) connection shared with NCAR from Boulder to the Denver Front Range GigaPoP (FRGP) via dark fiber. This link provides connections to the commercial Internet via AT&T, Cable & Wireless, and Level3 at 20 Mb/sec for each path. The NOAA/NCAR-FRGP link also provides a 310-

Mb/sec connection to Internet2. In addition, NOAA Boulder maintains a secondary 20-Mb/sec connection to the commercial Internet through the University of Colorado (CU). CU also maintains a separate 622-Mb/sec link to the FRGP that is available to NOAA Boulder should the NCAR link fail.

#### **PRTN**

PRTN maintains a 9 Mb/sec connection to the commercial Internet under its current Raytheon contract. In addition, the B-Site Forrestal Campus maintains a 100-Mb/sec microwave connection to the Main Campus of Princeton University. This connection provides access to Princeton University's 100 Mb/sec connection to Internet2.

#### **C.11.5.2 Proximity to additional wide-area networking capability**

#### **BLDR-1 & BLDR-2**

NOAA Boulder is a joint owner of the Boulder Research and Administrative Network (BRAN). BRAN is an eleven-mile dark fiber metropolitan-area network owned by NOAA, UCAR/NCAR, CU, and the City and County of Boulder. Besides interconnecting BRAN owner sites, it also has key tactical access to commercial communication providers, US West and ICG. NOAA has indisputable access to 24 strands of BRAN fiber. NOAA Boulder and the FRGP have contracted with the National Lambda Rail (NLR) network to attach in fiscal year 2005. Using Dense Wavelength Division Multiplexing, NLR will provide multiple channels of service at startup including: a 10-Gigabit Ethernet routed circuit to Internet and Internet2, multiple one-Gigabit Ethernet switched circuits to NLR nodes, multiple wavelengths for dedicated point-to-point circuits, and service for the Global Lambda Integration Facility.

#### **PRTN**

Level3 maintains a dark-fiber run along U.S. Route 1, which is roughly 1/3 mile from the lab across the field. The B-Site of Forrestal Campus is connected to the DOE-supported Princeton Plasma Physics Laboratory (PPPL) through dark fiber. PPPL is on DOE's ESnet wide-area network with a bandwidth capacity of at least OC3.

### **C.11.6 Fire Alarm And Fire Suppression Capabilities**

#### **C.11.6.1 Fire Alarm and Suppression Systems**

#### **BLDR-1 & BLDR-2**

Both computer rooms will have FM-200 fire suppression with VESDA fire detection systems. Both will have wet-pipe overhead sprinklers with 155°F trip point. CO<sub>2</sub> bottles are located in the room with FE-36 bottles surrounding the room. Alarms are tied to the building alarm system.

#### **PRTN**

The PRTN computer room currently is protected by a high-voltage fire alarm system and a wet-pipe water sprinkler system. However, this winter, these

systems will be replaced with the following: (1) a low-voltage alarm system with advanced alarm panel in the room adjacent to the computer room, (2) a dry-pipe water sprinkler system, (3) plumbing and other infrastructure to support future use of a gaseous fire suppression system (such as FM-200), and (4) a VESDA early smoke detection system.

#### **C.11.6.2 Fire Extinguishers**

##### **BLDR-1 & BLDR-2**

Fire extinguishers located within the computer rooms are CO<sub>2</sub> bottles. Fire extinguishers located within the hallways and computer room access areas are FE-36 gas bottles.

##### **PRTN**

Following the NCEP fire event in 2000, PRTN reviewed its fire procedures and eliminated any dry-chemical fire extinguishers from the entry paths to the Computer Room. The lab also met with the supporting fire departments, both Plainsboro and Princeton Plasma Physics Laboratory (PPPL) fire units, to explain the dangers of dry chemical extinguishers to computer equipment and to verify fire suppression procedures to be followed during fire emergencies.

#### **C.11.6.3 Response Time Of Local Fire Department**

Response time is defined as the time duration between fire alarm activation and the arrival of fire department personnel and equipment at the site.

##### **BLDR-1 & BLDR-2**

Boulder Fire Department provides fire-fighting capability. A station is located approximately one mile away and the response time is between 4 to 6 minutes.

##### **PRTN**

The Princeton Complex has measured the response time of the fire department (PPPL as first responders, backed up by Plainsboro Fire Department) to average 10 minutes (minimum of 7 minutes and maximum of 14 minutes), based on 6 events (2002-2004). These events were caused by unrehearsed false alarm events that were caused by faulty alarm sensors. [The Princeton Complex will be replacing the entire fire alarm system during the winter of 2004-05.]

### **C.11.7 Special Capabilities**

#### **C.11.7.1 Capabilities for Continuous Operation**

The following describes any capability that the facility may have to permit continued operation during a power outage or to ride through power anomalies.

##### **BLDR-1**

All electrical systems are fully backed up for uninterrupted service to cover for any length of power outage. The CRAC units are powered from emergency circuits, and will experience a shutdown until the EM generator comes online and transfers power (usually one to two minutes.) at which time they will automatically and sequentially restart. Once the EM generator starts, it will continue to operate

until clean and stable utility power is maintained for 30 minutes, at which time the generator will transfer the load back to utility power and shut down.

#### **BLDR-2**

No EM generator capacity is available for this facility; therefore, all equipment will be placed on to UPS power fed from utility. Should a lengthy outage occur, the UPS would provide uninterrupted service, until which time a graceful shutdown of computer equipment is required. CRAC units and cooling systems will remain on until all equipment is off and all static air has been forced from the compute nodes. The CRAC units will then shutdown, prior to exhausting the battery capability of the UPS systems.

The DSRC outage history is in the following table and shows one outage per year that would mandate a shutdown of the computer room.

#### Recent History of Power Loss Episodes at DSRC

FY2003	
DATE	DURATION
26 Mar 03	13 Minutes
13 May 03	13 Minutes
8 Sep 03	30 Seconds
11 Sep 03	2 Minutes
FY2004	
10 Jun 04	52 Minutes

#### **PRTN**

PRTN's facility environment is designed primarily for riding through a brief power outage, assuming the outage is sufficiently brief that the chillers will be able to automatically restart so that the temperature and humidity conditions in the computer room are maintained within acceptable limits.

The following table indicates environment related outages for the last three fiscal years. During these years, there were 29 documented power fluctuation episodes (14 in FY02, 11 in FY03, and 4 in FY04) in which the UPSs and PDUs maintained conditioned power to the computer systems and the chillers either continued to operate or else shutdown and then restarted automatically.

#### Recent History of Environment-Related Outages for PRTN

FY2002	
23 Jul 02	Mechanical (chiller) outage of 53-minute duration with 36 minutes of system recovery time
FY2003	
25 Feb 04	Power outage of 103-minute duration with 224 minutes of system recovery time

FY2004	
	No outages

#### **C.11.7.2 Duration Of Ride-Through At Full Load**

##### **BLDR-1**

Not Applicable – Maintain operation during power outage

##### **BLDR-2**

Designed for graceful shutdown during a power outage lasting longer than 10 minutes, with CRAC units shutting down last.

##### **PRTN**

The UPS systems for the current system provide a ride-through for the computer systems of 30 minutes on the 500-KVA UPS and 60 minutes on the 225-KVA UPS. However, the room air conditioning units and the chillers do not have generator support. If the power interruption is brief, Chillers #1 and 3 are likely to restart automatically. This recycling capability, combined with generator backup for the chilled water pumps, permits the facility to ride through brief power interruptions without the system going down.

#### **C.11.7.3 Capabilities for Remote Facility Management**

The following describes any capability that allows the facility to be monitored and managed remotely.

##### **BLDR-1 & BLDR-2**

BLDR-1 has SCADA (Supervisory Control and Data Acquisition) coverage for environmental monitoring. BLDR-2 is projected to have SCADA coverage as well. CRAC units also have building automation tie-in for alarms. Both systems have dial-out capability.

##### **PRTN**

Currently the computer room does not have remote monitoring capabilities. Such capabilities are being investigated for possible future implementation.

#### **C.11.7.4 Capabilities for Lights-Out Operation**

The following describes any capability that allows the facility to operate during off-hours without operators present.

##### **BLDR-1 & BLDR-2**

BLDR-1 currently has no operator coverage from 7pm to 7am every day. The SCADA environmental monitoring system will trigger an automated Emergency Power Off (EPO) if the temperature reaches 94oF.

##### **PRTN**

Currently the Computer Room does not have remote monitoring capabilities. Such capabilities are being investigated for possible future implementation.

**C.11.8 Physical security****C.11.8.1 Campus Security****BLDR-1 & BLDR-2**

Both facilities are located on federal property with federal police providing access to the grounds round the clock.

**PRTN**

Security personnel from the Princeton Plasma Physics Laboratory patrol the B-Site of the Forrestal Campus on a regular basis during off-hours. Plainsboro police are also available on call. Key sections of the exterior buildings are lighted at night. Security cameras observe these areas, digitally recording and displaying on monitors in the Operations Room on a round-the-clock basis.

**C.11.8.2 Security of Building Containing Computer Room****BLDR-1 & BLDR-2**

Federal police stationed at main entrance only during work hours. All entrances are monitored by video surveillance at all times.

**PRTN**

DOC Building Access Policies, involving NOAA ID cards and temporary badges, prescribe access controls for the Princeton Complex. NOAA issues these IDs according to prescribed procedures. During off-hours, access is limited to one building entrance. Building entrances are monitored from the Operations Room by a camera monitoring system with a digital recording system. The Operations staff carries out building security tours on a frequent schedule to monitor building environment, security, and facilities infrastructure.

**C.11.8.3 Computer Room Security****BLDR-1 & BLDR-2**

Computer room has restricted access provided by cipher locks at all entrances. A proximity badge and/or PIN number are required at all times.

**PRTN**

Computer Room access is controlled by a centrally managed cipher lock system to restrict room access to approved personnel. Sign-in procedures are required for individuals not on the approved personnel list. Cameras in the Computer Room allow Operators to monitor and digitally record areas of the room not viewable from the window of the Operations Room.

**C.11.8.4 Procedures For Contract Personnel To Access A Government Facility****BLDR-1 & BLDR-2 & PRTN**

All contract personnel must submit to a background check and will be issued an “affiliate” NOAA ID badge.

**C.11.9 Estimated Annual Facility Cost of Operation of Computer Room (Excluding Power Costs)****BLDR-1 & BLDR-2**

Annual costs for rent, maintenance, heating, and other general services were \$42.25/sq. ft. in FY 2004 as paid to the building landlord, GSA.

**PRTN**

Annual facility costs, excluding electrical utility costs, are estimated to be \$10/sq. ft. This includes: rent paid to Princeton under the Government triple-net lease; facilities services (physical plant maintenance, equipment service, janitorial services, etc.); utilities (other than electrical power); and wide-area network fees.

**C.11.10 Availability of blueprints****BLDR-1 & BLDR-2**

Blueprints are available from local GSA office and will need to be signed for by approved contractors with a need to review them.

**PRTN**

The Princeton Complex maintains a hardcopy archive of all blueprints for the Computer Building and associated chilled water plant upgrades.

**C.11.11 Projected Availability of Floor Space, Power, Cooling, and WAN Bandwidth**

The following amounts show projected total resources available at the indicated dates under the assumptions provided. The October 2004 availability is shown for reference purposes only, since Contractors cannot utilize the resources until October 2005.

		<b>Oct. 2004</b>	<b>Oct. 2005</b>	<b>Oct. 2006</b>
<b>Available Raised Floor Space (sq. ft.)</b>	BLDR	0	1,500	3,750
	PRTN	1,900	2,700	5,500
	<b>TOTAL</b>	<b>1,900</b>	<b>4,200</b>	<b>9,250</b>
<b>Available Non-Raised Floor Space (sq.ft.)</b>	BLDR	0	0	0
	PRTN	360	360	360
	<b>TOTAL</b>	<b>360</b>	<b>360</b>	<b>360</b>
<b>Available Power for New Hardware (kVA)</b>	BLDR	0	350	600
	PRTN	250	150	575
	<b>TOTAL</b>	<b>250</b>	<b>520</b>	<b>1,175</b>
<b>Available Cooling for New Hardware (tons)</b>	BLDR	0	146	218
	PRTN	180	155151	283275
	<b>TOTAL</b>	<b>180</b>	<b>301297</b>	<b>501493</b>

### Assumptions for Projections

#### Available Raised Floor Space:

**BLDR:**

- (1) October 2005 estimate assumes completion of computer room design, construction and testing.
- (2) October 2006 estimate reflects available space occupied by systems purchased under current HPTi contract. In reality, the new contractor will be required to coordinate new equipment installation with current system removal so as to minimize loss of compute cycles to the Government.

**PRTN:**

- (1) Oct. 2004 estimate assumes StorageTek silos have been moved onto non-raised floor as indicated in Fig. 3.
- (2) Oct. 2005 estimate assumes the mid-contract upgrade under the current Raytheon contract will provide 20% of additional free raised floor space compared to current floor-space usage.
- (3) Oct. 2006 estimate reflects available space occupied by systems leased under current Raytheon contract. In reality, the new contractor will be required to coordinate new equipment installation with current system removal so as to minimize loss of compute cycles to the Government.

#### Available Power:

**BLDR:**

- (1.) The projected available power is dictated by the cooling capacity that is available. Currently, the chilled water plant has reached its rated capacity and no further load can be supported without the addition of a new chiller.

**PRTN:**

- (1.) The projected available power assumes, somewhat arbitrarily, that total power to the computer room is limited to the current UPS capacities, which total 725 kVA, reduced by a 10% safety factor to roughly 650 kVA. This is a conservative estimate, given the total substation capacity, but does not include usage of chillers, which are currently run off the Main Building substation.
- (2.) Power estimates assume steady state ratings, not book value ratings for systems.
- (3.) Oct. 2005 estimate assumes a 25% increase in power usage over current usage.
- (4.) As with floor space, the Oct. 2006 reflects only minimal system usage from the current system with a background usage of 75 kVA.

Available Cooling:**BLDR:**

(1.) The chill water plant has reached its maximum capacity and will lose its N+1 capability on hot days. A failure within the chill water plant, on a hot day, will force the BLDR-2 computer room to load shed its cooling resources and shut down.

**PRTN:**

(1.) Assume the building load for the Princeton Complex for the hottest day of the year is 274 tons, based on observation that Chiller #1 alone (400 tons) will reach maximum capacity on this day when cooling the Princeton Complex plus the current system.

(2.) Compute cooling load for the system by multiplying system power usage (converted to tons) by 1.31.2 to reflect (conservatively) other room heat sources and cooling inefficiencies. (3.) Assume maximum chiller capacity uses chillers #3 (350 tons) and #4 (230 tons) operating together. Joint operation of chillers #1 (400 tons) and #4 (230 tons) are also viable within the current configuration. However, as indicated above, joint operation of chillers #1 and #3 will require electrical and mechanical (plumbing) modifications.

		Oct. 2004	Oct. 2005	Oct. 2006
<b>Available Wide-Area Network Capacity (Mb/sec.)</b> (list by type below)				
<b>Abilene / Internet 2</b>	BLDR	310	500	0
<b>Commercial Internet</b>		80	80	100
<b>National Lambda Rail (NLR)</b>				
<b>NOAA Research / Internet 2</b>		0	1,000	10,000
<b>Commercial Internet</b>		0	1,000	1,000
<b>*Note: 10,000 Mbps and/or dedicated Lambda's</b> <b>(wavelengths) available via NLR in 2005 and 2006</b> <b>if needed.</b>				
<b>Microwave to Internet2 (PU allowed)= 50 Mb/s</b>	PRTN	25	20	15
<b>Commercial Internet = 9 Mb/s</b>				
	<b>TOTAL</b>	<b>415</b>	<b>2,600</b>	<b>11,115</b>

### Assumptions for Projections

#### Available Wide-Area Network Capacity:

##### **BLDR:**

(1.) WAN capacity assumes timely completion of National Lambda Rail installation and connectivity

##### **PRTN:**

- (1.) Assume Princeton University only permits NOAA to use a total of 50% of its Internet2 bandwidth capacity.  
(2.) Assume current NOAA usage of total bandwidth will grow at a rate of 15% per year.

*[Figure 1 withheld from Public Website – Available upon request.]*

**Figure 1** *Projected Computer Room Layout for BLDR-1 in September 2006. The area enclosed within the dashed line indicates the floor space that will be available for Offeror use in October 2006. The small squares shown indicate 2'x2' floor tiles. Vent tiles are indicated by stippled squares.*

*[Figure 2 withheld from Public Website – Available upon request.]*

**Figure 2** *Schematic of Computer Room Layout for BLDR-2 facility that is expected to be available in October 2005*

*[Figure 3 withheld from Public Website – Available upon request.]*

**Figure 3** *Computer Room Layout for PRTN as of January 2005. Small squares in the figure indicate 2'x2' floor tiles. Vent tiles are indicated by stippled squares.*

*[Figure 4 withheld from Public Website – Available upon request.]*

**Figure 4** *Computer Building Schematic for Princeton Complex*

## Appendix C – Government Furnished Equipment (GFE)

The following describes Government-owned property that is available to be furnished to the Contractor for the performance of this contract. Because some Government-furnished equipment that is located at the designated sites is only available for use by the Contractor if that facility is used, the GFE lists are separated into two categories: “Site-Constrained GFE”, which is equipment that is only available if used at the designated site; and “Unrestricted GFE”, which is available for use under this solicitation without site restrictions.

### Boulder, CO

#### Site-Constrained GFE

Note: Partial use of this equipment is authorized up to the limits set forth in Appendix B.

Manufacturer	Part/Model Number	Description	Qty	Date Available
<b><i>Uninterruptible Power Supply</i></b>				
Chloride	EDP90/300/250 4x4	UPS Cabinet and Battery Bank (250kVA)	1	10/2006
	EDP70L/100/100 4x4	UPS Cabinet and Battery Bank (100kVA)	2	10/2006
<b><i>Computer Room Air Conditioning</i></b>				
Liebert	FH740C	Chill Water - Down Draft – 30 tons	4	10/2006
<b><i>Environmental Monitoring</i></b>				
Phonetics	SCADA 3000	Sensaphone SCADA 3000	1	10/2006

#### Unrestricted GFE

- None

### Washington, DC

#### Site-Constrained GFE

- None

#### Unrestricted GFE

- None

**Princeton, NJ****Site-Constrained GFE**

<b>Manufacturer</b>	<b>Part Number</b>	<b>Description</b>	<b>Qty</b>
<b><i>Computer Room Security Monitoring</i></b>			
		C-Cure 800	
Compac	4403US	Compac MS 2000 C-Cure 800 OS	1
Emerson	P761VBT0EENC	Monitor C-Cure 800	1
<b><i>Computer Room Security Monitoring</i></b>			
		Intellex 16000	
Sensormatic	DVMS DV16000	Intellex 16000 Video recorder system w/CD backup max 16 cameras	1
NEC	FE1250+BK	Monitor Intellex 16000	1
Pelco	CC-3700-S	Color CCD Camera	12
<b><i>Power Distribution Units</i></b>			
United Power	PDM4-F3-225-K13-426	PDU 1 (225 kVA)	1
United Power	PDM4-F3-225-K13-426	PDU 2 (225 kVA)	1
United Power	PDM4-F3-225-K13-426	PDU 3 (225 kVA)	1
<b><i>Uninterruptible Power Supply</i></b>			
MGE	72-131522-000	UPS Cabinet and Battery Bank 225 KVA	1
<b><i>Power Distribution Units</i></b>			
Liebert	PPA125C	PDU Unit (125 KVA)	1
EPE	PD084M48A12-125	PDU Unit (125 KVA)	1
<b><i>Computer Room Air Conditioners</i></b>			
DataFlow	CCT-60C4	CRAC 1 (35 Tons)	1
DataFlow	CCT-60C4	CRAC 2 (35 Tons)	1
DataFlow	CCT-60C4	CRAC 3 (35 Tons)	1
DataFlow	CCT-60C4	CRAC 4 (35 Tons)	1
DataFlow	CCT-60C4	CRAC 5 (35 Tons)	1
DataFlow	CCT-60C4	CRAC 7 (35 Tons)	1
APC	CCT-60C4	CRAC 8 (35 Tons)	1
APC	CM-3.0-W-BC-D	CRAC 6 (3 Ton)	1

**Unrestricted GFE**

<b>Manufacturer</b>	<b>Part Number</b>	<b>Description</b>	<b>Qty</b>
---------------------	--------------------	--------------------	------------

### ***Switch for Shared Storage***

SGI	FC-SWITCH-16	SGI Fibre Channel switch with 16 ports and one power supply	8
SGI	FC-SWITCH-PWR	Optional second power supply for SGI FC switches (8 & 16 port)	8
SGI	FC-SWKIT	Rackmount kit for mounting FC-SWITCH-8 or FC_SWITCH-16 in F1RACK	8
SGI	XSWOPTGBIC	Short Wave Optical GBIC kit containing 6 GBICs	16
SGI	XCOPGBIC	Copper GBIC kit containing 6 GBICs	6

### ***Switch for Tape Drives***

SGI	FC-SWITCH-16	SGI Fibre Channel switch with 16 ports and one power supply	8
SGI	FC-SWITCH-PWR	Optional second power supply for SGI FC switches (8 & 16 port)	8
SGI	FC-SWKIT	Rackmount kit for mounting FC-SWITCH-8 or FC_SWITCH-16 in F1RACK	8
SGI	XSWOPTGBIC	Short Wave Optical GBIC kit containing 6 GBICs	16
SGI	XCOPGBIC	Copper GBIC kit containing 6 GBICs	6

### ***Console and Monitoring***

SGI	SG230-00008	230L Workstation, 667 MHz PIII, 128MB PC 133 SDRAM, 20GB IDE, V3 Gfx 32MB DDR, Red Hat 6.1	2
SGI	91-AB945-001	19" Northern Hemisphere Monitor	2
SGI	91-AD001-001	Keyboard, Mouse, Speakers, Power Cords, Monitor Cable, User Manual	2
SGI	SSU60003	HW on-site support 4hr rresponse, 7x24, years 1-3 for SGI 230 Workstation	2
SGI	SC4-PCP-2.0	Performance Co-Pilot - Performance Monitoring tool for IRIX 5.3 and higher (replaces PCPORIGIN)	1
SGI	SC4-PCPHPC-1.0	Performance C-Pilot Add-On agent for IRIX 6.5 clusters (replaces SC4-PCPARRAY-1.0); requires SC4-PCP-2.0	8
SGI	SV4-PCPCOL-2.0-10	Performance Co-Pilot Collector 10 license pack	1

### ***Hierarchical Storage Mgmt System***

SGI	ORIGIN-3800	SGI Origin 3800 server - 64 CPUs (GFE) 600Mhz R14000A processors), 64GB memory, 16 Local Disk Channels (1Gb/sec), 16 Shared Disk Channels(2Gb/sec), 24 Tape I/O Channels, and 2 Gigabit Ethernet Channels. 180GB	1
StorageTek	9310002-0000	POWDERHORN 6000 CART/450 EPH	3

StorageTek	9940L03-0000	9940, Library, Fibre	24
StorageTek	9840L03-0000	9840, Library, Fibre	22
Sgi	FC-SWITCH-16	16 port 1Gb FC switch (8Cu/8Optical)	12
Sgi	ORIGIN-3800	Sgi Origin 3800 server - 64 CPUs (Lease 1a) 600Mhz R14000A processors), 64GB memory, 16 Local Disk Channels (1Gb/sec), 16 Shared Disk Channels(2Gb/sec), 24 Tape I/O Channels, and 2 Gigabit Ethernet Channels.	1

## Storage

***HSMS Local Storage***

Sgi TP9100 1GB	TP9100 D-Brick with fourteen 36GB 10KRPM Drives	16
STOR-CTRL 128	TP9100A Dual Channel Control Unit with 1Gb FC	16
Sgi TP9100 1GB	TP9100 D-Brick with fourteen 18GB 10KRPM Drives	2
STOR-CTRL 128	TP9100A Dual Channel Control Unit with 1Gb FC	2
Empty Rack	I/O Racks with AC Power Distribution	3
FC-SWITCH-16	16 port 1Gb FC switch (8Cu/8Optical)	4

***Hierarchical Storage Mgmt System***

StorageTek	9310002-0000	POWDERHORN 6000 CART/450 EPH	1
StorageTek	9840	Media Cartridges	3000
StorageTek	9940	Media Cartridges	12000

## Connectivity

***Connectivity - Existing***

Cisco	WS-C6509	Catalyst 6509 Chassis	1
Cisco	WS-CAC-1300W	Catalyst 6000 1300W AC Power Supply	1
Cisco	WS-CAC-1300W/2	Catalyst 6000 Second 1300W AC Power Supply	1
Cisco	SFC5K-SUP-5.5.1	Catalyst 6K Supervisor Flash Image, Release 5.5(1)	1
Cisco	WS-X6K-SUP1A-2GE	Catalyst 6000 Supervisor Engine 1A, Enhanced QoS, 2GE	2
Cisco	WS-X6416-GE-MT	Catalyst 56000 16-port Gig-Ethernet mod. MT-RJ	2
Cisco	WS-X6408A-GBIC	Catalyst 6000 8-port GE, Enhanced QoS	2

Cisco	WS-X6348-RJ-45	Catalyst 6000 48-port 10/100, Upgradable to Voice	2
Cisco	WS-G5486	1000BASE-LX/LH "long haul" GBIC	4
Cisco	WS-G5484	1000BASE-SX 'Short Wavelength' GBIC (Multimode only)	36
Cisco	CISCO3640	Cisco 3600 4-slot Modular Router-AC with IP software	2
Cisco	S364C-12.0.4T	IP	2
Cisco	NM-2CT1-CSU	2-Port Channelized T1/ISDN-PRI with CSU Network Module	2
Cisco	NM-1FE1CT1-CSU	1 Port F Ethernet 1 Port Channelized T1/ISDN-PRI with CSU NM	2

## Appendix D – Terms

### Abbreviations

ADIC	Advanced Digital Information Corporation
ANSI	American National Standards Institute
CM2	GFDL's Climate Model Version 2
CONOPS	DOC's Concept of Operations acquisition process
COTS	Commercial Off-The-Shelf
CXFS	SGI's shared filesystem for storage area networks (SANs)
DDN	Direct Data Networks
dGB	Decimal Gigabyte (10 <sup>9</sup> bytes)
dPB	Decimal Petabyte (10 <sup>15</sup> bytes)
DOC	US Department of Commerce <a href="http://www.commerce.gov">http://www.commerce.gov</a>
DSRC	David Skaggs Research Center
ENTA	Enterprise Network Target Architecture
FC	Fiber Channel
FISMA	Federal Information Security Management Act of 2002
FSL	NOAA's Forecast Systems Laboratory <a href="http://www.fsl.noaa.gov">http://www.fsl.noaa.gov</a> 325 Broadway Boulder, CO 80305
FYxxxx	Fiscal Year xxxx
Gb	Gigabit (2 <sup>30</sup> bits)
Gb/s	Gigabit(s) per second
GB	Gigabyte (2 <sup>30</sup> bytes)
GFDL	NOAA's Geophysical Fluid Dynamics Laboratory <a href="http://www.gfdl.noaa.gov">http://www.gfdl.noaa.gov</a> Princeton University Forrestal Campus 201 Forrestal Road Princeton, NJ 08542
GFE	Government Furnished Equipment
GNU	Unix-compatible software system
GPFS	General Parallel File System
GrADS	Grid Analysis and Display System
GSA	General Services Administration
HBA	Host Bus Adapter
HFS	Home File System
HIMF	Hallberg Isopycnal Model (Fortran)
HPC	High Performance Computing
HPCS	High Performance Computing System (see definition)
HPSS	High Performance Storage System
HSI	Hierarchical Storage Interface

HSM	Hierarchical Storage Management
HSMS	Hierarchical Storage Management System
htar	HPSS tape archiver
IDL	Interactive Data Language
IEEE	Institute of Electrical and Electronics Engineers
IMSL	International Mathematical and Statistical Library
I/O	Input/Output
IPCC	Intergovernmental Panel on Climate Change
IT	Information Technology
ITS	Information Technology Services
KB	Kilobyte (2 <sup>10</sup> bytes)
KVA	Kilovolt amperes (1000-volt amps).
LAN	Local Area Network
LPAR	Logical Partitioning
LSC	Large-Scale Computing
LTD	Live Test Demonstration
Mb	Megabit (2 <sup>20</sup> bits)
Mb/s	Megabit(s) per second
MB	Megabyte (2 <sup>20</sup> bytes)
MPI-2	Message Passing Interface (latest version)
MTBF	Mean-Time-Between-Failure
NAG	Numerical Algorithms Group
NCAR	National Center for Atmospheric Research
NCEP	NOAA's National Centers for Environmental Prediction <a href="http://www.ncep.noaa.gov">http://www.ncep.noaa.gov</a> 5200 Auth Road Camp Spring, MD 20746
NetCDF	Network Common Data Format
NIST	National Institute of Standards and Technology
NLR	National Lambda Rail
NOAA	DOC's National Oceanic and Atmospheric Administration <a href="http://www.noaa.gov">http://www.noaa.gov</a>
NWS	National Weather Service
OAR	Office of Oceanic and Atmospheric Research
OCIO	Office of the Chief Information Officer
OMB	Office of Management and Budget
OPM	Office of Personnel Management
PB	Petabyte (2 <sup>50</sup> bytes)
PDU	power distribution units
PFTP	Program to transfer data from host to host
PM	Preventive Maintenance
POSIX	Portable Operating System Interface
PSE&G	Public Service Electric and Gas
R&D	Research and Development
RAID	Redundant Arrays of Inexpensive Disks

RFI	Request For Information
RFP	Request For Proposal
RUC	Rapid Update Cycle
SAN	Storage Area Network
SGE	Sun Grid Engine
STK	Storage Technology
TB	Terabyte (2 <sup>40</sup> bytes)
UDUNITS	Unidata's Units Library
UPS	Uninterruptible power supplies
WAN	Wide Area Network
WRF	Weather and Research Forecast
WRF-CHEM	Weather and Research Forecast with Atmospheric Chemistry
WRF-SI	Weather and Research Forecast with Static Initialization

### C.11.12 Definitions

Application Memory	The maximum resident set size used by an application process in Mega/GigaBytes.
Availability	The availability level of a computer, component, or device is a percentage figure determined by dividing the operational use time by the difference between wallclock and null time.
Communication Fabric	The hardware component(s) supporting MPI message traffic.
Community Supported Software	Open Source Software
Degraded Mode	System operation at less than normal capability due to the loss of hardware or software components on that system.
Downtime	That period of time when all of an HPCS component's workload cannot be accomplished due to a malfunction in the Contractor-maintained HPCS hardware or software, or when the HPCS or a component of the HPCS is released to the Contractor for maintenance.
HPCS	A High Performance Computing System can be one large System or an aggregation of Subsystems. A given System/Subsystem can be further described by individual components. A single component may span multiple Subsystems.
Null Time	The period of time when the workload cannot be accomplished due to environmental failure, such as loss of electric power or cooling, or recovery from environmental failure.

Operational Use Time	The time during which equipment is available to the Government, exclusive of preventative maintenance time, remedial maintenance time, or Contractor-caused machine failure. Partial credit may be given by the Government for equipment operating in degraded mode (for example, when a portion of the processors, memory, disk, etc. on a computer is unavailable). The Government may declare the entire HPCS down even if parts of the HPCS are available.
Physical Memory per Processor Core	The size of the memory chipset in Mega/GigaBytes supporting a single processor core.
Preventative Maintenance (PM)	That maintenance performed by the Contractor which is designed to keep the equipment in proper operating condition. It is performed on a scheduled basis.
Processor Core	The component of the processor containing the dependently scheduled floating point and integer registers and arithmetic and load/store unit(s).
Processor Socket	The motherboard component designed to receive the processor.
Remedial Maintenance (RM)	That maintenance performed by the Contractor which results from Contractor-supplied equipment or operating software failure. It is performed as required and therefore on an unscheduled basis.
Suite	A set of concurrent instances of a workstream and possibly, a number of given workstreams (see target IT architecture).
Subsystem	
Target IT Architecture	The set of hardware and software systems which accomplish a workstream suite.
Total (Workstream) Throughput Time	The wallclock time from submission of the first instance of a workstream component to the successful program end of the last instance of a workstream component. In general, multiple workstream types and instances may be running on a particular target IT architecture. The Total Throughput Time for a workstream is defined in the context of all concurrent loads for the target IT architecture.
Workstream	A single instance of end-to-end processing.