Statement of Need (SON)

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 a 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 the 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.

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. A workstream (WS) is a single instance of end-to-end processing (including pre- and post-processing).

The Government requires a single Contractor to be responsible for the design, installation, maintenance, and support of the HPCS. 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 is comprised by Subsystems that are designed to help accomplish NOAA's R&D missions, as represented by workstreams. Some of the components that comprise a Subsystem include: Large-Scale Computing, Development Computing, 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, and potentially be 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, the R&D workstreams are comprised of computationally intensive environmental modeling applications coupled to I/O-intensive codes and extensive 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.

Of the annual funding, 94% will be dedicated to the components and services for the HPCS. Under direction from the Government, the remaining 6% of the funds will be reserved 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 amount of the reserve funds may be unilaterally increased by the Government at any time during the course of the contract. The Government is not under obligation to direct the 6% reserve 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 solely at the Government's discretion.

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 provide balanced performance.

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 an increase over current capabilities in computational throughput for NOAA R&D. At least one significant 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 provisions of this SON establish the Government's minimum acceptable capabilities of the HPCS based on NOAA's experience in performing its mission. However, innovation in proposed high-performance solutions is encouraged. Newer technologies or an approach different from that presented here may provide opportunities to further 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 to 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, requiring considerable computational resources. These models are developed through research and development (R&D) efforts within NOAA that occur primarily at 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) use models for study of atmosphere, ocean, climate, air quality and ecosystem behavior.

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 of 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 and greater then 15 million files. See Appendix A for a more complete description of the current HPCS.

C.4 Benchmarks

C.4.1 Workstream (WS) Profile

As previously mentioned, 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 three different computer systems. 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 frequent intervals during the contract. Further, the Government would 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 performance upgrade during both the base and option periods. It is desirable that the mid-life performance upgrades for the base period for WS7 – WS9 occur between Q4FY2007 and Q1FY2008 while WS1 – WS6 performance upgrades should occur between Q2 and Q3FY2008.

Benchmarking for system upgrades shall be handled as follows:

- Benchmark revision is required at the *option period upgrade* (see Section C.9.1 for information on the option period upgrade).

- Benchmarks may change by mutual agreement between Government and Contractor at *mid-life upgrades*.
    - 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.
    - Further, any software changes that are allowed will be used to create a new baseline with respect to the existing system.
    - Thus, a performance increment is always defined with respect to a constant set of source code across the change.

- Revised workstream benchmark codes and performance baseline data must be delivered to the Contractor no later than 6 months prior to delivery of an upgrade.

- Throughput and Scaling performance data shall be reported as described in Section J:
    - Throughput and Scaling values with class A changes only.
    - Any additional Throughput and Scaling values with class B, C and D changes. All such changes will be evaluated for risk and approved in the same manner as the RFP response.

- 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, this 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. These levels indicate current estimates of NOAA support for the programs represented by the workstreams. Offerors are expected to use this information in establishing performance goals for the benchmark. The offered configuration should provide flexibility to the Government to meet changing priorities throughout the contract period. The funding begins for workstreams 7-9 in FY2006. The funding for the remaining workstreams will begin in FY2007. The Government requires 6% of the cumulative annual "Total" funding amount (the far right-hand column) cited in Table I below be reserved. Based on information provided by the Contractor, the Government will evaluate proposals to verify that workstream performance is appropriate to the 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 Sections C.9.2 and C.9.3.

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	$0.7	*$ 3.0*
	FY2007	$4.0	$6.0	$4.0	$1.933	$1.933	$1.934	$1.3	$1.0	$0.7	*$ 22.8*
	FY2008	$4.0	$6.0	$4.0	$1.933	$1.933	$1.934	$1.3	$1.0	$0.7	*$ 22.8*
	FY2009	$4.0	$6.0	$4.0	$1.933	$1.933	$1.934	$1.3	$1.0	$0.7	*$ 22.8*
Option	FY2010	$4.0	$6.0	$4.0	$1.933	$1.933	$1.934	$1.3	$1.0	$0.7	*$ 22.8*
	FY2011	$4.0	$6.0	$4.0	$1.933	$1.933	$1.934	$1.3	$1.0	$0.7	*$ 22.8*
	FY2012	$4.0	$6.0	$4.0	$1.933	$1.933	$1.934	$1.3	$1.0	$0.7	*$ 22.8*
	FY2013	$4.0	$6.0	$4.0	$1.933	$1.933	$1.934	$1.3	$1.0	$0.7	*$ 22.8*
Total		*$28.0*	*$42.0*	*$28.0*	*$13.531*	*$13.531*	*$13.538*	*$10.4*	*$ 8.0*	*$ 5.6*	*$162.6*

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 whenever 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 and level as the failed processors. The capability of the LSC to operate in degraded mode during repairs 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 Subsystem 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 Interactive Component

For each LSC Subsystem, binary compatible interactive development resources shall be provided. The interactive 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 by the Resource Management Software (see Section C.5.3.1) without a system reboot.

The Government desires that the operation of 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 to interactive work, and vice versa, without a reboot or a restart to the batch system.

The Government desires early access to the interactive resources for training, software porting and tuning, and OS and application configuration testing. Early access to initial system delivery and subsequent system upgrades (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 targeted for that subsystem 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 this access 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.5.2.8).

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 HSMS shall meet the performance requirements of the HSMS Archive benchmark given in Section J.1.4.4.

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 minimum of a two-tiered storage scheme for its data archive, comprised of frequently accessed nearline storage (disk storage or tape storage robotically mounted at high speed) and offline storage (tape storage robotically mounted). This 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 desires that files recalled from the offline tier automatically revert to nearline tier residency. User-specified migration between tiers is also desirable. The Government desires a file display command, comparable to UNIX “ls –l”, which shows each file’s residency in the nearline or offline tier. The Government requires a daily report, submitted to the COTR, of storage used by groups and users in all storage tiers.

If disk storage is used for the nearline tier it must be fault-tolerant and sized for the Government’s file size distribution and data access patterns. As it may be difficult to recover from disk storage failure, the Government requires all HSMS disk files to either have a secondary copy or a backup within 24 hours.

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

NOAA’s existing data archive for workstreams 4-6 resides on Government-owned StorageTek Powderhorn silos located at IBM’s facility in Gaithersburg, MD (2 silos) and the Government facility in Fairmont, West Virginia (1 silo) (see Appendix A). These silos presently support NOAA’s Operational Central Computer System

(OCCS) and its backup. These legacy HSMS silos are available to the Contractor as unrestricted GFE beginning in FY2007 (see Appendix C). The Contractor shall make arrangements to keep this data available. The Government desires a transparent migration of this legacy archive data to the new HSMS, maintaining the present filesystem image. The HSMS proposed to support workstreams 4-6 is required to interface with either the OCCS or its backup which are located at Fairmont, W.V. and Gaithersburg, MD. (See C.5.2.5)

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 September 9, 2006. At that point, the HSMS is available to the Contractor as unrestricted GFE (see Appendix C). The Government desires a transparent migration of this legacy archive to the new HSMS, maintaining the present filesystem image. The Contractor shall make arrangements to keep this data available during the transition, potentially at a reduced level of performance.

Multiple archives may be provided. Each data archive shall be presented to the associated subsystem as a single filesystem. Initially, per Subsystem, the HSMS shall be able to support 200 million files and individual files that are up to 1 TB in size. The Government desires HSMS software comparable to the San Diego Supercomputer Center (SDSC) Storage Resource Broker (SRB), which transparently combines logically related small files into one larger "container" file in the archive filesystem. More information on SRB is available at http://www.npaci.edu/DICE/SRB.

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. As directed by the Government, within 30 calendar days, the Contractor shall provide a media recovery service to recover data in the event of nearline or offline storage media failure.

For the tape storage, 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. The "average search time" must apply to the method the HSMS software uses to position 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 tape storage, the Government desires several types of tape media optimized for the file size distribution, balancing file positioning time with file read time.

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 Interfaces to the NOAA Operational Central Computer System (OCCS) and the Backup OCCS

The exchange of information to and from the NOAA Operations HPC is an integral part of NOAA's overall Research, Development and Operations processes; an objective of this procurement is to preserve and enhance the ability to execute such exchange of information.

The data that is generated from the OCCS, and available at both Fairmont, WV, and Gaithersburg, MD, is required to be written to the R&D HSMS that supports workstreams 4-6. This data consists of the run histories from the operational models along with observation data that are outputs from the OCCS. Table IIa shows the daily volumes that NOAA expects to be written over the life of the contract. Current practice is to "batch" these data and write them to the R&D HSMS during off-peak hours to avoid contention with other HSMS users.

The OCCS backup system has a requirement to be able to both read from and write to the R&D HSMS that supports workstreams 4-6. The OCCS backup system will require similar HSMS access and latency characteristics as those that are described in Section J.1.4.4.3 for workstreams 4-6. This data is composed of model run histories and data generated from pre-operational model development work. Table IIa shows the daily volumes that NOAA expects to be written and read over the life of the contract.

C.5.2.6 Data Generation Profile

C.5.2.6.1 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 = Workstream 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.5.2.6.2 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*
Base	FY2006				2.4	2.2	1.6	*6.2*
	FY2007	2	3	2	2.4	2.2	1.6	*13.2*
	FY2008	4	3	4	4.8	4.4	3.2	*23.4*
	FY2009	5	3	5	4.8	4.4	3.2	*25.4*
Option	FY2010	8	6	8	9.5	8	6	*45.5*
	FY2011	10	6	10	9.5	8	6	*49.5*
	FY2012	17	6	17	12	10	8	*70*
	FY2013	24	12	24	12	10	8	*90*
Option FY2014		24	12	24	12	10	8	*90*

C.5.2.6.3 Data Generation for OCCS and the Backup OCCS

Table IIa shows the projected data generation for the OCCS and the Backup OCCS.

Table IIa -- Data Generation / Access Profile for OCCS and Backup

	OCCS (Write)	Backup (Read)	Backup (Write)
FY2005	1.0 TB/day	0.8 TB/day	0.8 TB/day
FY2006	1.6 TB/day	1.2 TB/day	1.3 TB/day
FY2007	2.2 TB/day	1.7 TB/day	1.8 TB/day
FY2008	3.2 TB/day	2.5 TB/day	2.7 TB/day
FY2009	4.6 TB/day	3.5 TB/day	3.9 TB/day
FY2010	6.6 TB/day	5.1 TB/day	5.5 TB/day
FY2011	9.4 TB/day	7.2 TB/day	7.9 TB/day
FY2012	13.0 TB/day	10.0 TB/day	11.0 TB/day
FY2013	19.0 TB/day	14.0 TB/day	16.0 TB/day
FY2014	26.0 TB/day	20.0 TB/day	22.0 TB/day

Note: The totals shown in the OCCS (Write) and Backup (Write) columns of this table are in addition to the totals shown in Table II (Data Generation for WS4-WS9).

C.5.2.7 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% for the first year; 50% for the second year; and 15% as a persistent archive.
WS5	100% for the first year; 50% for the second year; and 15% as a persistent archive.
WS6	100% for the first year; 50% for the second year; and 15% as a persistent archive.
WS7	100% for 3 DAYS; 20% for the first year; 10% for 3 years; and 5% for 8 years.
WS8	100% for 3 DAYS; 20% for the first year; 10% for 3 years; and 5% for 8 years.
WS9	100% for 3 DAYS; 20% for the first year; 10% for 3 years; and 5% for 8 years.

The data retention profile for the data shown in table IIa is 100% for the first year; 50% for the second year; 30% as a persistent archive. Persistent archive is defined as the data in the archive that will exist or remain indefinitely. The funding for this storage is contained in the funding profile for workstreams 4, 5, and 6 shown in table I.

The following figure shows the accumulated data retention for WS4 – WS9. This is utilizing the data generation numbers shown in Table II of C.5.2.6 and Table III of this section. WS1 – WS3 cannot be shown without the computational growth being determined.

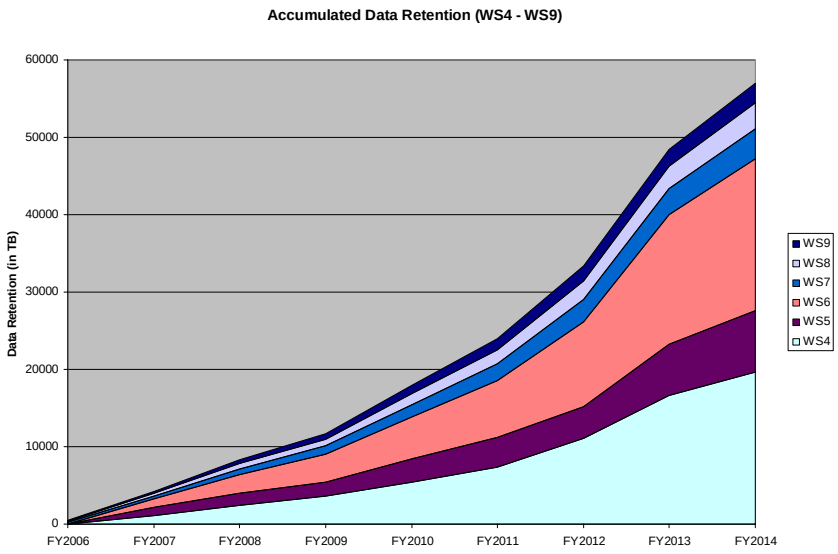

Figure 1 - Data Retention for WS4 - WS9

C.5.2.8 Automated backup

Automated backup shall be provided for all of the following: 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 made 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 file systems for which the backups are made.

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 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. All Resource Management Software shall be provided with documentation, training, and an established problem resolution process.

The Government requires resource management software that provides:

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

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 Batch Queuing Software to manage all HPC Subsystem computing resources. The Batch Queuing Software shall be provided with documentation, training, and an established problem resolution 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.
- Provide 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 users, groups, and projects or accounts.
- For accounting purposes, tag each job with a project or account string.
- Collect cluster-global job accounting information for each job, including:
    - the user, group, and project or account
    - the number of processors dedicated to the job
    - the job-total cpu time
    - the elapsed (wallclock) time
    - the date and time of job completion
    - NOAA specified programs identifier

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 an operator specified time window

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.
- 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, which 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 site licenses are provided.

- Matlab (http://www.mathworks.com)
- Mathematica (http://www.mathematica.com)
- IDL (http://www.rsinc.com/idl/)
- S-Plus (http://www.insightful.com/products/splus/)
- 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)
- Environment Modules (http://modules.sourceforge.net/)
- 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.tcsh.org)

Required application software for interactive use is comprised of:

- Ferret (http://ferret.wrc.noaa.gov/Ferret/)
- GEMPAK/N-AWIPS (http://my.unidata.ucar.edu/content/software/gempak/)
- 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 ncview 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 ncview 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 those involved with weather workstreams such as WS 4, 6 and 7-9 generally use 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 as well as either or both interactive and 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. 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 75% of those users are located at the David Skaggs Research Center (DSRC), in Boulder, CO. The remaining 25% 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 unstructured type using NCAR graphics (NCL and user applications built with NCAR graphics library), Ncview, IDL (interactively and applications), GemPAK, Vis5D, and others. Some of these data-browsing activities involve viewing single frame images while others utilize loops of fields over time as well as interactive 3D manipulations. These visualization tools are X-window based applications producing 24bit RGB images. Typical resolutions range between 800x600 and 1280x1024. At the time of the initial delivery of resources supporting WS7-9, a minimum of 40 simultaneous data browsing sessions producing graphics files to be animated by GFE workstations on the LAN in Boulder, CO shall be supported. The Government does not presently expect significant growth in the number of simultaneous users**.** The files produced will be up to 150 frames and should be delivered to the user's GFE workstation in less than two minutes after the generation of the graphics files is completed. Additionally, a minimum of 6 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. The files produced will consist of up to 150 frames and should be delivered to the user's GFE workstation in less than two minutes after generation of the graphics files is completed. 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% that are 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. The archive is currently mounted to each user workstation over the LAN. Thus, ncview may, run as a local application on the workstation, transferring the data from the file sitting on the archive staging disk. Some users currently prefer to copy the whole data file over the LAN to local storage on their workstation, again running ncview on the user workstation. Still others run ncview from a session logged onto the Analysis cluster (see Section C.10.2.3) directly in their /archive directory, thus sending a series of RGB frames over the LAN. 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. The Contractor is free to offer compression technology compatible with the user toolset which might reduce WAN/LAN requirements.

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. Note that within the time frame of initial delivery, the LAN supporting the Princeton, NJ user set is anticipated to be upgraded to gigabit Ethernet. Additionally, a minimum of 10 simultaneous interactive TotalView debugger sessions shall be supported with LAN-like interactive use for the Princeton, NJ user set. For information concerning the current SAN and LAN supporting these activities, see Sections C.10.3.2.3 and C.10.4.2.3, respectively.

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.5.2.6 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. The data required for workstreams 7, 8, and 9 have a 75% intersection with 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 will be negotiated.

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 is acceptable. At the discretion of the Government, further delays of the data can be counted as downtime. Also, it is required that no 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 should be noted that the data necessary for WS 4-6 is and will be available in the Washington D.C. metro area. It should also be noted that the data necessary for Workstreams 7-9 is and will be available in the Skaggs building in Boulder. It is uncertain if the necessary bandwidth will be available to GFDL in Princeton. If the workstreams are targeted for computational platforms not co-located with their current data

sources, the Contractor is required to provide any necessary bandwidth, above and beyond the bandwidth provided by the Government, for communication of this data.

C.5.4.4 Communications Requirement to Support the HSMS

The contractor shall be responsible for the communications link between the HSMS at either the Fairmont, WV or Gaithersburg, MD data access points, that is necessary to support Workstreams 4-5-6 and the Primary & Backup OCCS write and read functions described in table IIa. The interface specifications shall be negotiated with NCEP.

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) and Department of Commerce (DOC). 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, the US DOC IT Security Program Policy and Minimum Implementation Standards, 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 (C&A) package. All IT equipment delivered as part of this procurement is to be considered a Government computing resource, regardless of its location or actual owner. The Government will initiate the C&A package for this system and then transition maintenance of the security documentation to the Contractor after delivery and acceptance. The Contractor may also elect to employ equipment which is not part of the delivery for use in remotely diagnosing, monitoring, or managing the delivered system. If this Contractor-owned equipment has any special access or trust relationship with the delivered system, then the Contractor must initiate and maintain any appropriate C&A documentation for this equipment.

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.

While the Government may provide some measure of physical or network perimeter protection for the system, the Contractor shall be responsible for IT Security of the delivered solution. In instances where communication or interaction is required across system boundaries, standard (NIST SP 800-47) interconnection agreements will define the security responsibilities and procedures for all parties involved.

All solutions shall comply with all DOC and NOAA security policies. 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"). However, the final system categorization will be determined through development and approval of the System Security Plan (SSP).

The following Commerce Acquisition Regulation (CAR) security clause is applicable to this contract:

Security Processing for Contractor/Subcontractor Personnel Working on a Department of Commerce Site (Low and Moderate Risk Contracts) (CAR 1352.237-72 (MAR 2000)

Security Processing Requirements

(1) U.S. Citizens Working on DoC Site

All contractor (and subcontractor) personnel proposed to work on the premises of a Department of Commerce site for 180 days or more must undergo security processing by the Department's Office of Security (OSY) to be eligible to work on the site.

(2) Foreign Nationals (Non-U.S. Citizens)

Regardless of anticipated length of on-site work, all foreign nationals to be employed under this contract must:

Have legal visa status with the Immigration and Naturalization Service (INS);
Have advance approval from the serving Security Officer in consultation with the Office of Security.

b. Submittal Requirements – U.S. Citizens

(1) Duration of Onsite Work: 180 to 364 days (between 6 months and 1 year)

For individuals who will be performing work on a DoC site between 180 and 364 days, the Department will perform a Special Agreement Check (SAC). The scope of the SAC will include checks of the Security/Suitability Investigations Index (SII), other agency files (INVA), Defense Clearance Investigations Index (DCII), FBI Fingerprint (FBIF), and the FBI Information Management Division (FBIN).

The contractor must complete and submit the following form to the Contracting Officer's Technical Representative (COTR):

Form FD-258 (Fingerprint Chart)

Copies of this Form can be obtained from the COTR. Upon receipt of the FD-25, the COTR will complete form OFI 86C (Special Agreement Check) and will forward both to the operating unit Security Officer. The Security Officer will advise the COTR whether work can commence prior to suitability determination, based on the specifics of the situation. The COTR will notify the Contractor of an approved contract start date as well as favorable findings of the suitability determination.

(2) Duration of Onsite Work: 365 days (1 year) or more

Individuals proposed to perform work on a DoC site for 1 year (365 days) or more are required to have a NACI check (National Agency Check Plus Written Inquires).

The Contractor must complete and submit the following forms to the Contracting Officer's Technical Representative (COTR):

Standard Form 85P (SF-85P, Questionnaire for Public Trust Positions), and FD-258 (Fingerprint Chart).

Copies of these Forms can be obtained from the COTR. Upon receipt of the required forms, the COTR will forward the forms to the operating unit Security Officer. The Security Officer will advise the COTR whether work can commence prior to suitability determination, based on the specifics of the situation. The COTR will notify the Contractor of an approved contract start date as well as favorable findings of the suitability determination

c. Submittal Requirements – Foreign Nationals

All Foreign nationals proposed to work on a DoC site will be subject to a Special Agreement Check (SAC) to determine whether the foreign national has official legal status in the United States.

The Contractor must submit the following forms to the COTR for all foreign nationals proposed to work on a DoC site:

FD-258 (Fingerprint Chart)
Form OFI 86C (Special Agreement Check) with signature authorization for release of information

Copies of these Forms can be obtained from the COTR. Upon receipt of the required forms, the COTR will forward the forms to the operating unit Security Officer. The COTR will notify the Contractor of favorable findings and will notify the Contractor regarding an approved date to commence work under the contract.

d. Suitability Updates

Any individual (including foreign nationals) processed on the form OFI-86C (Special Agreement Check) who stays on the contract over 364 days will be required to have a NACI complete suitability check to stay on the job site.

e. Notification of Disqualifying Information

IF OSY receives disqualifying information on a contract employee, the Contractor, upon notification of such by the Contracting Officer, must immediately remove the employee from duties which require access to DoC facilities.

Individuals may be barred form working on the premises of a facility for any of the following:

(1) Conviction of a felony or a crime of violence or of a misdemeanor involving moral turpitude.

(2) Falsification of information entered in security screening forms or on other documents submitted to the Department.

(3) Improper conduct once performing on the contract, including criminal, infamous, dishonest, immoral, or notoriously disgraceful conduct or other conduct prejudicial to the Government regardless of whether the conduct directly related to the contract.

(4) Any behavior judged to pose a potential threat to departmental personnel or property.

Failure to comply with the requirements may result in termination of the contract. Compliance with these requirements shall not be construed as providing a contract employee clearance to have access to classified information.

C.5.6 Facilities Requirements

The Contractor shall provide a facilities proposal for housing the proposed equipment to meet the Government's requirements over the life of the contract. Contractors may use any of the following to meet the Government's needs: Any of the available Government facilities (as described in Section 0), either as is or with modifications; Contractor-provided facilities; or a combination of Government facilities and Contractor facilities. The Contractor shall provide details of their proposed facilities solution in the Facilities Proposal that they include in their submission.

C.5.6.1 Government-Provided Facilities

Contractors may use the projected facilities offered at the NOAA locations in Boulder, CO, Princeton, NJ, Fairmont, WV, and Largo, MD, which are described in Section 0. Section C.11.11 indicates the Government-supported resources that are projected to be available, beginning at the times indicated and 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 in their proposal that indicates proposed changes, with schedules, for the work to be done. The modification plan must be in accordance with site restrictions.

Contractor use of the Government-supported resources assumes the following conditions apply:

- The Government is responsible for delivering the indicated facilities resources (although the Contractor will be responsible for paying power utility costs in accordance with Section H.19). Because of this, 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 the funding of any facility modifications (e.g., new electrical equipment, plumbing, additional air conditioning equipment) needed to install and support the proposed systems, and to obtain facility resources in excess of what the Government makes available as indicated in Section C.11.11. The Contractor will also be responsible for carrying out the facility modifications at facilities such as the Princeton facility where the contractor is permitted to perform the build-out; for these facilities, the Contractor must include a description of these modifications, including schedules, as part of its Facility Proposal.

- Any new facility equipment (e.g., air handlers, chillers, UPSs, or switchgear) or building modifications will become the property of the owner of the facility, with the exception of UPS equipment at PRTN. The new facility equipment will be maintained by the owner of the facility, unless otherwise negotiated with the Contractor. However, UPS equipment at PRTN, either GFE or vendor-owned, that is used to support systems under this contract shall be maintained by the Contractor. All other UPS equipment will be maintained by the Government or the facility owner.

- For those facilities, such as the GSA-managed facilities, in which the any facility modifications must be paid to the facility manager by the Government prior to the work being contracted for and performed by the facility manager (e.g., GSA), the Government shall be fully reimbursed by the Contractor for the work to be performed through credits against the contract that do not extend beyond the fiscal year in which the expenditures are made by the Government.

For any facility modifications to be made, the Contractor must follow the following procedures that are specific to each of the indicated sites:

Boulder, CO
The David Skaggs Research Center (DSRC) is owned by the Government and operated by the General Services Administration (GSA). Because of this, all building modifications are subject to GSA Regulations, thereby requiring that GSA contractors carry out these modifications. 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.

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. The Contractor is advised that, because of the long implementation delays for large projects (see table below), facility modifications, at least during the first year of the contract, should be limited to small projects such as power receptacle changes. Large facility modifications will require substantial lead-time to complete

Examples of Boulder Facility Modifications (Estimations ONLY)

Small Facility Modifications		
Modification	**Estimated Cost[1]**	**Estimated Time Required**
Power Receptacle Relocation (per receptacle)	$1500	60 Days
Power Receptacle Alteration (per receptacle)	$2000	60-90 Days
Large Facility Modifications		
Additional Chiller and Cooling Tower	$1.2 Million	8 to 10 Months
Additional Emergency Generator	$1 Million	8 to 10 Months
[1] Note: Costs include installation, support equipment and associated building modifications. Costs do not include environmental studies, feasibility studies, building impact surveys or GSA overhead and program management fees.		

All facility infrastructure modifications and equipment at DSRC will become the property of the government and will be fully maintained by the government in accordance with manufacturers' maintenance schedules. Normal Government contracting procedures can be expected for any modifications.

Princeton, NJ

Because the Princeton buildings are owned by Princeton University, all construction shall be performed and paid for by the Contractor and is subject to all applicable building codes and approvals for the work from Plainsboro Township and Princeton University. In addition, if these facility modifications cause any changes to the exterior of the building, these modifications shall be in accordance with Princeton University's Princeton Forrestal Center Design and Development Criteria, which includes approval by the Design Review Committee.

Fairmont, WV

The NASA Independent Verification and Validation (IV&V) Facility is owned and operated by the West Virginia University Research Corporation (WVURC) but provided by WVURC to NASA to house NASA's IV&V Facility. All construction is subject to all applicable building codes and approvals for the work by WVURC and NASA. Any proposed modifications or upgrades will be approved by and installed by West Virginia University Research Corporation. NOAA will provide NASA/ WVURC with the requirements and funding for any facility modifications to be made and will be reimbursed by the Contractor through credits against the contract; NASA/WVURC will then have the modifications performed.

The following pricing from NASA/WVURC is provided for the cost to Offerors for possible facility upgrades to the Fairmont facility.

Electric Power Upgrades:

- Switchgear/UPS upgrade to provide circuits for additional PDUs $1,000,000
- Additional diesel generator (required for N+1 capability) $600,000
- Additional Power Distribution Units, cabling, labor $300,000

Cooling Upgrades:

- For a system using <300KVA, additional cooling units $200,000
- For a system using >300KVA, also need an additional chiller for N+1 $600,000

The time required to complete the above actions, once NASA/WVURC receives the requirements, is estimated to be 8-10 months.

Largo, MD

The Largo facility is a commercial facility that is leased through GSA. Modifications to the Largo facility are routed through the NOAA Facilities office, which contacts GSA and handles lease modifications. GSA contacts the building owner, Michael Management Inc., which in turn brings in contractors and handles any required permits. Modifications to date have had turnaround times of 2 to 6 months for completion (10 months for the generator, due to manufacturer delays), dependent upon scope. GSA personnel have consistently monitored ongoing work and pushed to meet the Government's time constraints.

C.5.6.2 Contractor-Provided Facilities

The Contractor may propose to use facilities other than the Government-provided facilities that are described in Section 11, Appendix B. Any Contractor-provided facilities shall meet 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, Section C.12, but are responsible for all shipping costs, and for the shipment both to the Contractor site and back as directed by the Government.
- Designated Government personnel must have access to the Contractor facility on an as-needed basis, subject to reasonable controls.

C.5.6.3 Office space, workstations, LAN connectivity, phone service, and office equipment for on-site contract personnel

The following policies will be in force for any Contract personnel to be located at Government sites.

The Government will provide LAN connectivity to all Contractor personnel located at its sites. The Contractor shall provide workstations, phone service, and office furniture for all of its personnel located at Government sites, although the Government may, for its convenience, choose to provide this on a case-by-case basis. To assure that IT security policies are maintained, the Government will require Contractors to install, operate, and maintain their workstations on the LAN in a manner that is consistent with Government IT security policies.

Boulder, CO
The Government will supply two (2) 150-sq. ft. offices and approximately 100 square feet of lab space for use by the Contractor.

Washington, DC
NCEP can accommodate two on-site personnel at the World Weather Building and its annexes. However, office space is extremely limited, so contractors should limit office space requirements as much as possible.

Princeton, NJ
Effective October 2006, the Government will provide two offices for use by Contractor's personnel who are responsible for system operation. These two offices will be occupied by the current Raytheon system support staff until the end of the contract on September 30, 2006. They are the large room adjacent to the Operators Room and an office across the hall. If the R&D Contractor requires additional office space and/or storage space beginning in October 2006 or requires office and/or storage space prior to October 1, 2006, it may construct

offices on the hard pan in the southwest corner of the Computer Room as part of its facility proposal. It may construct these rooms in any portion of the hard pan space that was previously planned for offices as indicated in Figure 3. If the Contractor does not choose to use this hard pan for this purpose, it may use this non-raised floor for installation of equipment that does not require raised floor.

In addition to the two offices provided for personnel responsible for system operation, the Government will attempt to provide additional office space to other Contractor personnel on an as-needed and as-available basis.

Fairmont, WV
The Government will provide approximately 500 sq. ft. of office space to be configured as directed by NOAA.

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 (as defined in Section 0, Appendix D) will be used in the determination of the actual System Life Throughput. Periods of Remedial and Preventive Maintenance count as downtime. Null time will not be counted as downtime. Null time is that period of time when the workload cannot be accomplished due to environmental failure at a Government-furnished site, such as loss of electric power or cooling, or recovery from environmental failure. 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 ultimately determined 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 software. (Routine procedures for entry to a facility shall not be construed as denial of access.)

A component's downtime shall end when it 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.

The Government, at its discretion, may provide a series of carefully monitored jobstreams that may include up to twenty-five (25) individual batch jobs in order to verify the reliability of the system. These jobtreams will be monitored for end-to-end success. Complete or partial failure of any intermediate step will result in declaring the jobstream to have failed. Success ratio for a jobstream will be determined as the number of successful jobstreams divided by the number of attempted jobstreams in a 30-day period per subsystem. If multiple types of jobstreams are being monitored on a subsystem, then the aggregate of their success ratios will be computed by weighting by the total amount of CPU-hours each jobstream should use over a 30-day period. The Government requires that ratio be greater than or equal to the proposed availability. This will be evaluated monthly. Failures due to null time, application errors, or data errors will not be included in the numerator or denominator of the above ratio. The Government is responsible for ensuring that the applications function properly on the subsystem. The Government may, at its discretion, choose to accept tools, support, and advice that will make the jobstreams more reliable.

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 with associated output data 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 targeted to run on that system. For example, if Subsystem A typically runs workstreams 1 – 2, workstreams 3 – 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, excluding 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 = 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).

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

$$SLT_{sample} = 11446.4 throughput_benchmark$$

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. At the sole discretion of the Government, shortfalls in throughput, when compared to the proposed throughput, shall be made up with either the delivery of additional equipment installed at no cost to the Government, or accrual of downtime credit. Should the Government elect the delivery of additional equipment to satisfy throughput shortfalls, the Government will calculate the duration of the compensation using the demonstrated benchmark performance on the upgraded HPC Subsystem. 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. Should the Government elect downtime credits in lieu of additional equipment, downtime credits will be deducted from the most recent invoice. Downtime credits shall accrue on all components of the system(s).

At the sole discretion of the Government, shortfalls in proposed availability or performance of non-computational components shall be rectified with the delivery of additional equipment or engineering, at no cost to the Government, or accrual of downtime credit. 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 throughout the life of the HPCS. If the parties are unable to reach mutual agreement, the Government reserves unilateral right to determine changes to the HPCS workstream benchmarks and the Contractor agrees that it will comply subject to the Disputes Clause.

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 each of 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.5.6.3 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 the Contractor to maintain a current, itemized list of all Contractor-supplied hardware and software items in printable electronic form. The listing must be updated whenever a system upgrade or engineering change proposal results in a change to the system configuration or delivery of additional equipment. The listing is to be submitted to the COTR upon request.

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" requires increased integration of NOAA 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:
    a) achieve code interoperability with good performance
    b) provide continuity of operations
    c) respond to changing programmatic requirements
    d) 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 that some elements may compromise performance unacceptably. In that case, NOAA may choose to forego or delay 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.

C.8.4 Configuration and Change Management Plan

The Contractor shall establish a Configuration Management (CM) process including a Configuration Management Control Board to include both contractor and NOAA Government 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 at the Annual Meeting 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.8.5 Transition Requirements

C.8.5.1 The Contractor shall develop and submit, within 60 days of a written request by the Government, a comprehensive Transition Baseline Report. The goal of the Transition Baseline Report is to avoid disruption of the day-to-day conduct of HPC computing while achieving a smooth and orderly transfer of responsibility from the Contractor to a successor contractor. It should compile available information related to current system configuration, performance, support requirements, issues, risks, and other information deemed pertinent by the Government. This plan shall be submitted to the Contracting Officer and COTR.

C.8.5.2 Thirty (30) calendar days prior to the completion of this contract, an observation period shall occur, during which time personnel of the successor contractor may observe operations and performance methods of the outgoing Contractor. This will allow for orderly turnover of facilities, equipment, and records and will help to ensure continuity of service. The outgoing Contractor shall not defer any requirements for the purpose of avoiding responsibility for or of transferring such responsibility to the succeeding contractor. The outgoing Contractor shall fully cooperate with the succeeding contractor and the Government so as not to interfere with a smooth transition and ongoing accomplishment of the Government's work.

C.8.5.3 Thirty (30) calendar days prior to a scheduled midlife system upgrade the contractor shall submit a written plan to the Government describing how the upgrade will be implemented. It should contain a schedule, information related to configuration, performance, support requirements, issues, risks, and other information deemed pertinent by the Government. This plan shall be submitted to the Contracting Officer and COTR.

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 estimates that the annual funding for this extension 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 estimates that the annual funding for this extension 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 order to accommodate these, and similar situations, the contract includes, as an option, system components and workstreams to augment the R&D HPCS as necessary. The optional system augmentations are comprised of several major elements to include workstreams (CLINS 0010A through 0010J inclusive), compute systems (CLIN 0010K), storage systems (CLIN 0010L), network (CLIN 0010M and 0010R), software (CLIN 0010N), workstations (CLIN 0010P), and maintenance (CLIN 0010S) as identified in Section B.

If additional workstream(s) should be acquired under this option, it is understood and agreed that the actual configuration acquired may be different than the workstream configurations identified in Section C of the contract. In such case(s), the actual configuration being acquired will be negotiated.

Furthermore, in addition to the items identified in Section B of this contract, the Government may, during the term of this contract, require goods and services not otherwise specifically identified, but within the scope of this contract. Should requirements arise, these additional goods and services will be negotiated under the Changes clause of this contract.

C.9.5 Engineering Support

The Government requires an option to acquire 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

A senior applications analyst 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.

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, but 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 Contractor 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.

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

C.9.5.4 Visualization/Data Engineer

NOAA may require professional engineering services in support of its HPC facilities. These services may be for Visualization, Graphics, or Data Engineering. Activities may include, but are not limited to the following:

* Generation and manipulation of graphical products.
* Data flow optimization.
* Creating and using tools for accessing, manipulating, viewing and representing the data.
* Data Conversion.
* Metadata manipulation for downstream post-processing.
* Data format evolution and porting.
* End user support.

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 times. 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 MB/s 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 GB 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 Gaithersburg, MD has access to two StorageTek Powderhorn silos, each with eight 9940B tape drives and about 1.25 PB capacity. The two silos are fully populated with tape cartridges. NOAA's HPCS in Fairmont, WV has access to one StorageTek Powderhorn silo, with four 9940B tape drives and about 1.25 PB potential capacity. The silo currently holds about 2300 tape cartridges. NCEP's experience indicates, on average, each 9940B cartridge can hold 250 GB. The HSMS is presently considered part of the operational system at NCEP. This storage system 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 IHSI Unix-like interface. Additional information on HPSS and IHSI can be found at:

http://www.hpss-collaboration.org
http://www.sdsc.edu/Storage/hsi/

NCEP currently writes about 2TB/day to tape and by the fourth quarter of 2005 all available capacity will have been consumed. The tape archive layout is the same at both the Gaithersburg and Fairmont locations; the storage devices are connected to the computational cluster by dual Gigabit Ethernet connections through a Cisco 6513 router. Each computational cluster supports "interactive" nodes where users control HPSS functions.

C.10.3.2.3 Princeton, NJ

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.

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.

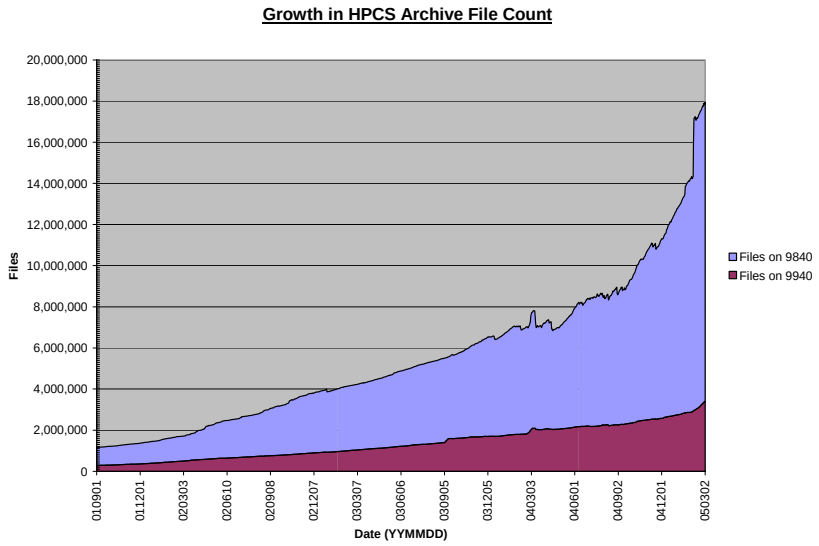

Figure 2 - Growth in tape-resident files on HSMS in Princeton

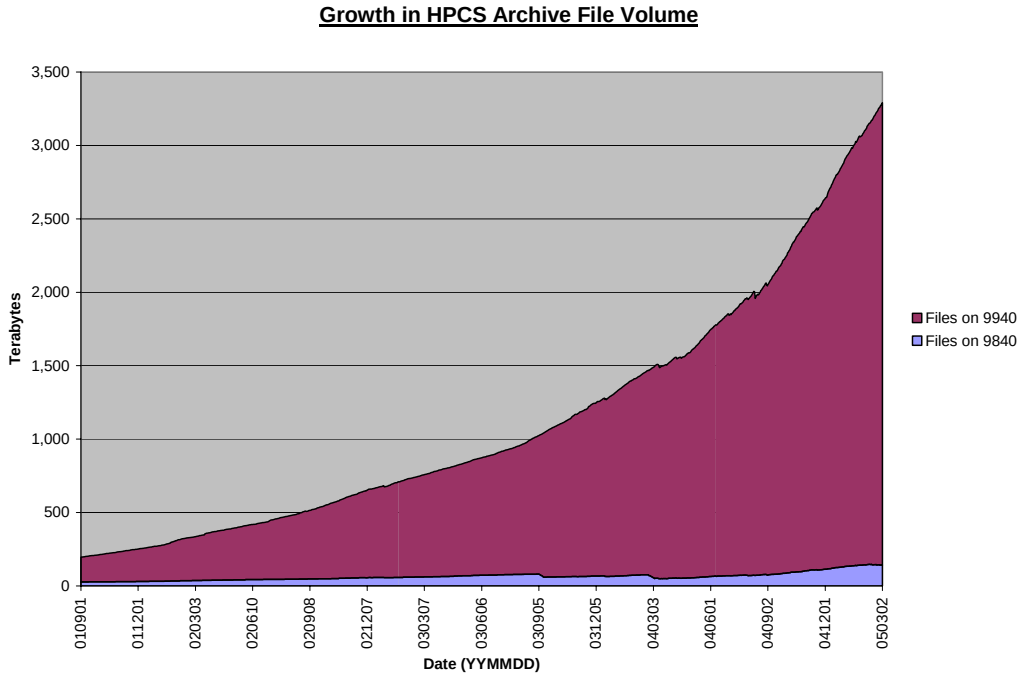

Figure 3 - Growth in tape resident storage on HSMS in Princeton

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	15 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. This setup is inherited 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 switch 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.

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 6 is in use on the primary set of clusters. 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 (DELETED). 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 to 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 can be found 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 Data Flow Diagrams

C.10.6.1 Boulder, CO

The following two diagrams show the data flows for the RUC2 backup model and WRF model that are run at the Forecast Systems Laboratory

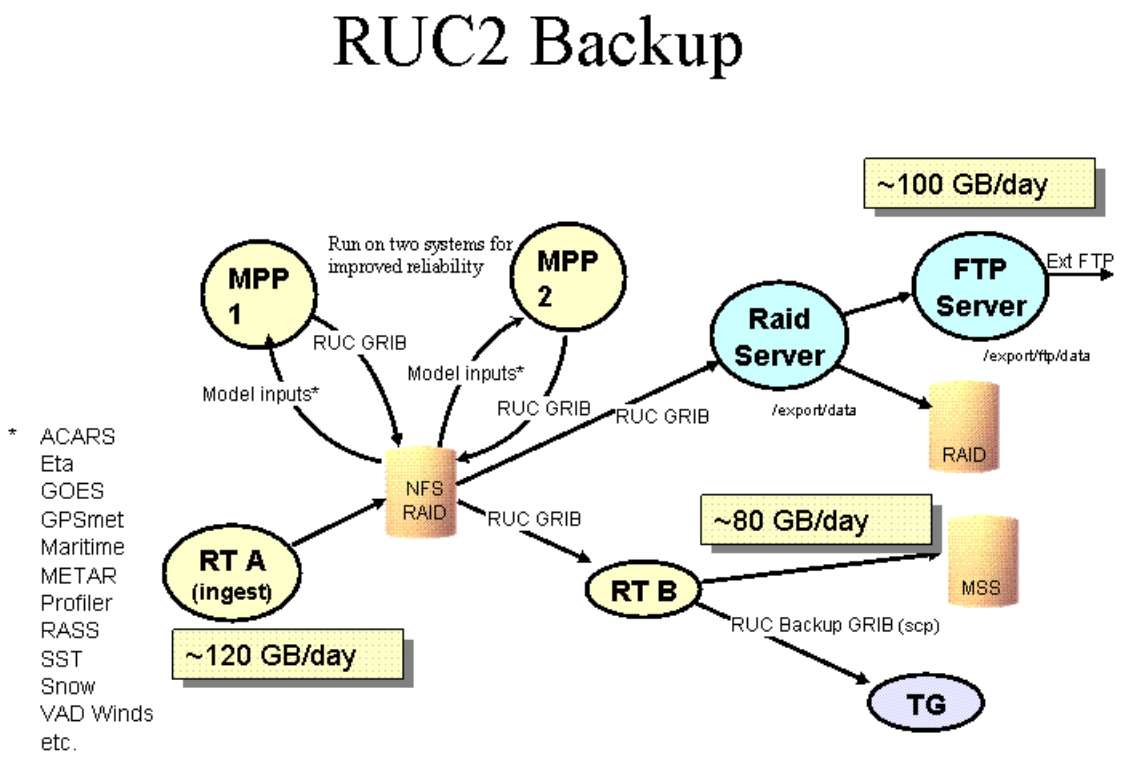

Figure 4 – RUC2 Backup

WRF

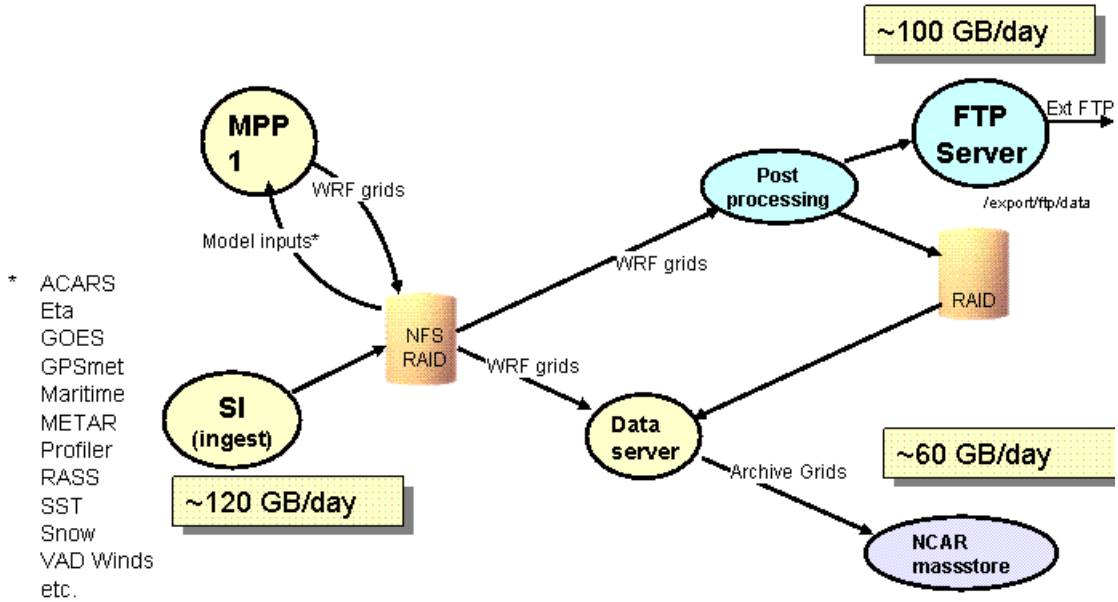

Figure 5 – WRF

C.10.6.2 Washington, DC

The following diagrams show the current data flows and target data flows for workstreams 4-6.

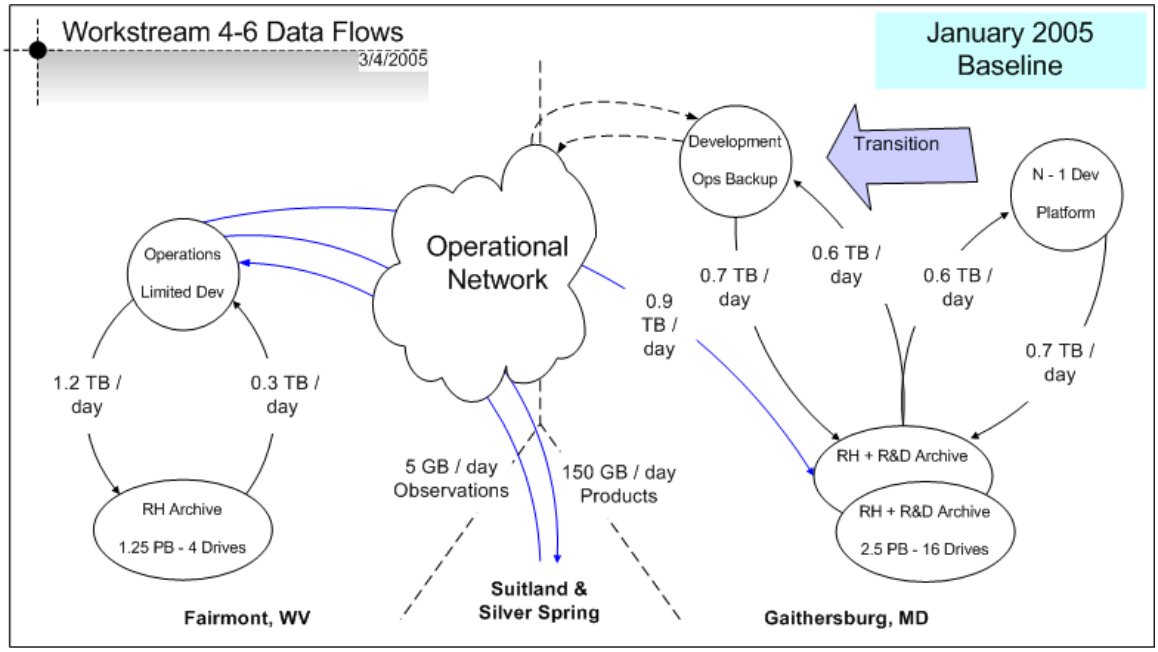

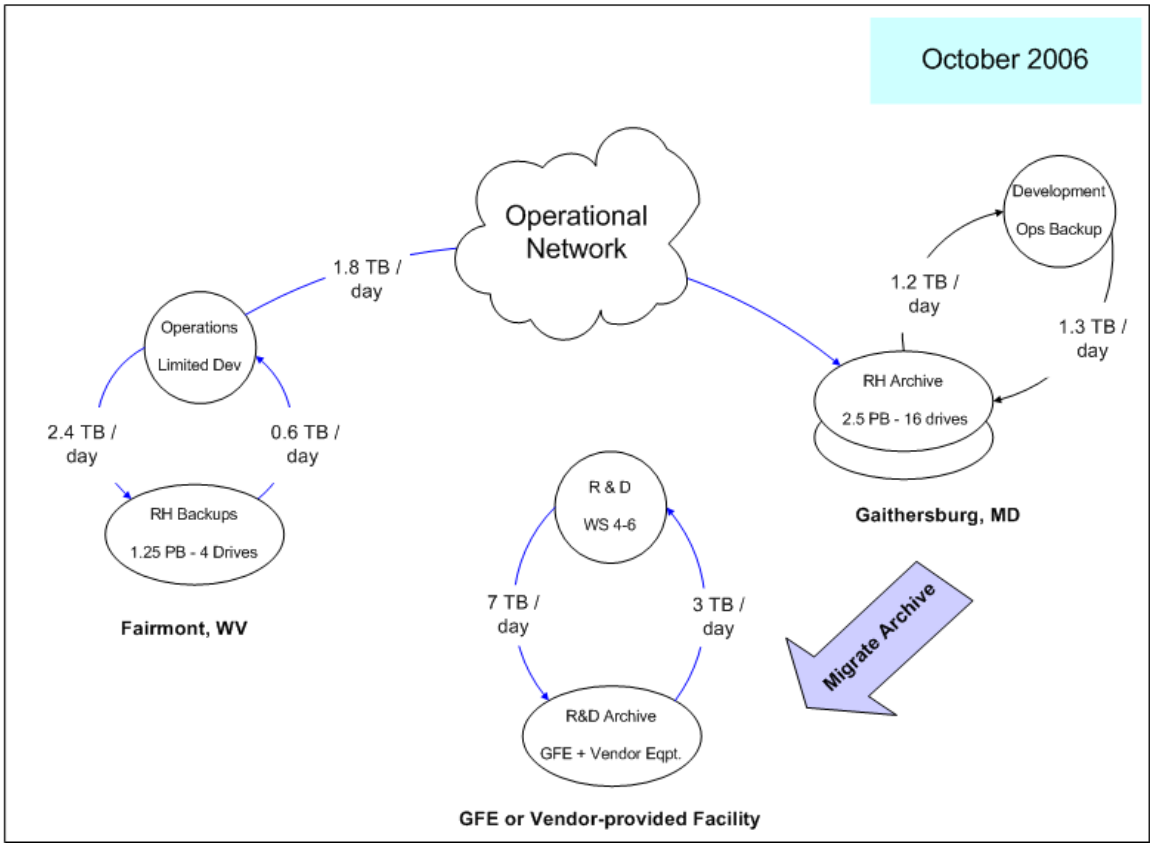

Figure 6 – Workstreams 4-6

C.10.6.3 Princeton, NJ

The following data flow diagram depicts a single climate run.

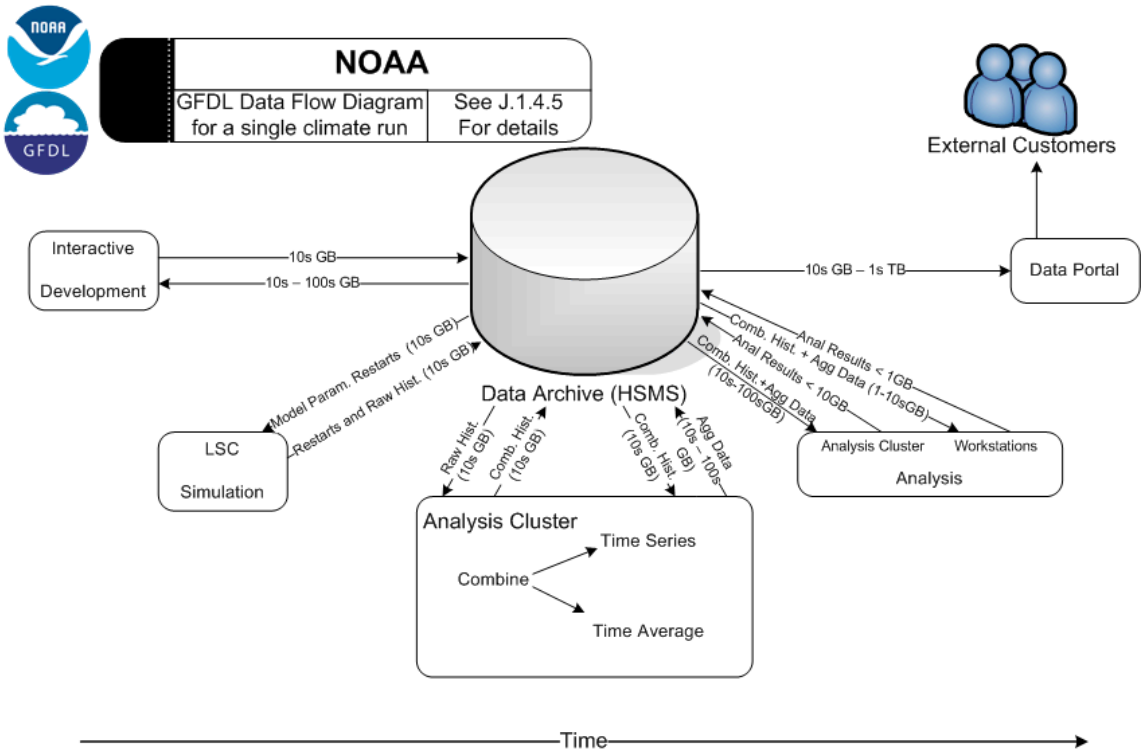

Figure 7 – Single Climate Run

C.11 Appendix B - Available Government Facilities

This document provides descriptions of five government-furnished facilities, four of which are designed to support high-performance computer systems and the fifth which involves an existing computer room and office space within NOAA's administrative computing facility. All five facilities are available for possible use by Offerors to house systems under the NOAA HPC R&D contract, including the availability of facility resources as projected in Section C.11.11 under the assumptions provided therein and subject to the facility terms and conditions provided in Section C.5.6**.** Two of these facilities, designated as BLDR-1 and BLDR-2, are located in the David Skaggs Research Center (DSRC), 325 Broadway, Boulder, CO. DSRC is owned by the Government and managed by the General Services Administration (GSA). 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. The Princeton Complex is owned by Princeton University and leased to the Government. The fourth facility, designated as FAIRMONT, is space located in the NASA IV&V Facility which is leased from NASA by NOAA to house NOAA's backup NCEP supercomputer. The term of the current lease runs through September 30, 2008, but is expected to be continued. The fifth facility, designated as LARGO, is located in a commercial facility, known as the NOAA Information Technology Center (ITC), that is leased by NOAA through GSA for housing NOAA's administrative computer systems. The facility's address is ISMO, 1221-D Caraway Court, Largo, MD 20774-5381. The NOAA ITC is managed by the Information Systems Management Office (ISMO) of NOAA's Financial and Administrative Computing Division. The current lease of the facility containing LARGO is a five-year lease that began in mid-February 2004 and will end one day earlier than this date in 2009. The lease contains an option for a five-year extension beyond the facility's base period. At this juncture, it appears there are no plans to remove Commerce Administrative Management System (CAMS) from the NOAA ITC. Per GSA, lease cost is the lowest available within near proximity to the Capital Beltway. The Largo Metro Station is opening December 18 literally at the back fence, which will enhance access to the facility.

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. FAIRMONT currently house NOAA's operational backup supercomputer. The LARGO facility currently contains some of NOAA's administrative computer systems but has additional computer room space available to support components of the R&D HPCS. This space will be available for use in March 2005.

C.11.1 Layout and Physical Dimensions of Computer Room

C.11.1.1 Computer Room Layout

Facility Drawings
Offerors requesting facility drawings and related documents that contain sensitive but unclassified information 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.

Figures 1, 2, 2a, 3, 3a, and 6 referenced below will be provided to Offerors upon request and are not subject to these GSA regulations. Figures 4 and 5 contain sensitive but unclassified information and are subject to the GSA regulations; they will be provided to Offerors from whom the Government has received acceptable non-disclosure forms for the facilities shown in the figures (i.e., PRTN and LARGO).

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.

Figure 2a shows a revised room layout for BLDR-2 that seeks to address concerns expressed by Contractors in their questions as well as design requirements. The following changes have been made in the design:

- CRAC unit positioning changed (opposing)
- Secondary egress installed
- XDO units moved to accommodate racks that may be up to 48 inches deep
- Full 4 foot “cold aisle” between the fronts of the HPC racks.

PRTN
Figure 3 shows the computer room layout for the current system, but with the following planned modification: move silos onto the northern 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 shows the overall Computer Building layout. The rooms adjacent to the Computer Room as 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.
- Office for Contractor support personnel
- 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.

FAIRMONT

Figure 6 shows the space NOAA has leased from NASA within the NASA IV&V Facility. NOAA is currently leasing 6,800 sq. ft. of space, but this space could be expanded to include an additional 500 sq. ft. to provide office space in support of this contract.

LARGO

Figure 5 indicates the rooms that are available within the NOAA ITC. The current Xerox Room involves 600 sq. ft. of raised floor with dimensions of 20 feet by 30 feet. The Media Room has 320 sq. ft. of raised floor with dimensions of 16 feet by 20 feet.

Figure 6 shows the layout of the surrounding rooms adjacent to the offered facility.

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.

FAIRMONT
The Fairmont facility is located within a technology park adjacent to Interstate 79 running North and South through West Virginia. The following website can provide addition information regarding directions and local information: http://www.ivv.nasa.gov.

LARGO
Space to be made available involves the Xerox Room (20 ft x 30 ft) and the Media Library (16 ft x 20 ft). Other areas could be made available through relocation of the Lab/Conditioned Storage area and/or converting an office area, which was originally computer-conditioned space. These last two areas do not currently have raised floor, so installation of raised floor would involve additional cost. Specifically, one concrete floor area with a 12-foot-high ceiling could be converted to raised floor at either 724 sq. ft. by retaining the existing standalone office, or 864 sq. ft. with the office removed. An adjoining concrete floor area of 576 sq. ft. could be made available and could be added as contiguous space, though removal of a wall, removal of existing unused plumbing, and raising the ceiling would be required.

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 2368 sq. ft. A portion of this space is used for non-HPC equipment and safety egress corridors. 1424 sq. ft. of raised floor space is available for equipment. An adjacent room contains an additional 156 sq. ft. of raised floor space, and is available as an assembly area.

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.

FAIRMONT
Figure 6 shows the current raised floor space at the NASA IV&V Facility that is leased by NOAA for NOAA's backup NCEP supercomputer. The room is a large open area (68 ft. X 96 ft.) with support pillars distributed throughout the room. An additional area (10 ft. X 37 ft.) is also leased by NOAA for office/administrative

use. Also, an additional 500 sq. ft. could be made available by NASA for this proposal if needed. NOAA is currently only occupying approximately half of the large area for the backup NCEP supercomputer thus leaving approximately 3,200 sq. ft. of the large room available for this proposal with an optional 500 sq. ft. of office/administrative space.

LARGO
Figure 5 indicates the available rooms. The current Xerox Room involves 600 sq. ft. of raised floor with dimensions of 20 feet by 30 feet. The Media Room has 320 sq. ft. of raised floor with dimensions of 16 feet by 20 feet. Possible additional raised floor space could be obtained in the amount of either 724 or 864 sq. ft., as discussed above in C.11.1.2. Thus the raised floor space available without modification totals 820 sq. ft. Installation of additional raised floor could produce a total of up to 1,784 sq. ft. An adjoining space of 576 sq. ft. could also be available, currently as non-raised floor space.

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 1424 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.

FAIRMONT
As indicated above, there is approximately 3,200 sq. ft. of raised floor space currently available at the Fairmont facility with an optional additional 500 sq. ft. available for office/administrative space. Even though the space is currently available, modifications to meet any necessary power and cooling requirements may not be completed until the summer of 2006, depending upon such requirements.

LARGO
As indicated above, there is a total of 920 sq. ft. of raised floor in the largo facility. This space will be available in March 2005. Additional raised floor space can be added to this facility through installation of more raised floor by the Contractor, producing a total of up to 1,784 sq. ft.

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. Roughly 360 sq. ft. of this non-raised floor, located adjacent to the doors leading to the loading dock, will be available as a staging area for new equipment. 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 spring to the locations on the hardpan as indicated in the figure.

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

LARGO
There is non-raised floor space available in the facility that can be used for storage and assembly as well as office space.

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 currently provided on the left hardpan in the Computer Room and in the UPS Room. Two of the offices will be available to the Contractor's system support personnel beginning in October 2006. The Contractor may use the non-raised floor space in the southwest corner of the left hard pan to construct additional offices, use as storage space, and/or use for locating equipment that does not require raised floor. Prior to the end of the current Raytheon contract, half of this space (32’x18’) will be reserved for Raytheon use. The offices that were planned to be constructed (as indicated in Figure 3) on this hard pan space are now offered to the Contractor to use to construct offices and/or storage space or else for installing equipment that does not need to be located on raised floors. In addition, the larger contractor room located on the loading dock side of the Operator Room will be left as is for contractor use as office space (available to the R&D contractor after the current Raytheon contract ends) and not divided into two offices as previously planned.

FAIRMONT
As stated earlier, approximately 500 sq. ft. of additional raised floor could be provide by NASA to accommodate office/administrative space or could be used as needed by NOAA or the contractor.

LARGO
Some of the existing space in the facility can be re-configured to provide office space for contractor support personnel. With no modifications to existing office area, the facility could accommodate up to 10 work spaces. With modifications, 30 or more office space may be available.

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

FAIRMONT
Drop ceiling height = 10 feet

LARGO
The ceiling height above raised floor space is 9 feet. In the concrete floor 724/864 sq. ft. area it is 12 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 84”. A waiver of this restriction may be allowed, up to 94 inches, depending on the location and placement of the equipment. Modifications to the overhead sprinklers, with approval from the GSA Fire Protection Engineer, may be necessary. It is strongly recommended that vendors who are contemplating the use of tall equipment submit proposed drawings of equipment layouts and specifications to the government prior to submission of their proposal. The government will attempt to provide a rough estimate of the costs that will be born by the vendor to accommodate the equipment.

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

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

LARGO
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).

FAIRMONT
Raised floor height = 24 inches

LARGO
Raised floor height = 18 inches

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 about 50’ to the computer room
BLDR-2: 25’ from the Loading Dock to the freight elevator and then about 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.

FAIRMONT
Loading Dock - Standard, designed to accommodate 18-wheeler semi-trucks. There is a powered overhead door with manual locks that is tied into the security system. The loading dock/staging area is approximately 600 sq. ft. The pathway into the raised floor area from the loading dock area is accommodated with double door ways that are 9 ft. high and 7 ft. wide. The pathway also runs over raised floor; thus, additional floor protection must be used for transporting racks over 1,000 lbs. over the raised floor

LARGO
Loading Dock - Standard, designed to accommodate 18-wheeler semi-trucks. There is a powered overhead door with manual locks that is tied into the security system. The loading dock/staging area is 600 sq. ft. The path into the raised floor area is from the loading dock up a reinforced ramp. The ramp area and path to it are 6 feet 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.

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

LARGO
Equipment may be staged and assembled in areas outside of the computer room.

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^2; 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^2; 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./ft²; Ultimate Load: 3850 lbs.; Rolling Load: 1000 lbs. (10 Passes)

Raised floor Specs (adjacent hallways): Concrete floors

PRTN
Computer Room Raised Floor Specs: Nevamar tile and CEI Bolted Stinger: Static Load: 1250 lbs; Ultimate load: 2900 lbs. Rolling load: 1200 lbs. (10 passes).

Most of the raised floor, indicated as 2' x 2' square tiles in Figure 3, was replaced to the above specs 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.

FAIRMONT
Computer Room Raised Floor Specs: Tec-Cor II Screwed Stringer: Concentrated Load: 1250 lbs <0.095"; Uniform Load: 300 lbs < 0.040"; Ultimate Load: 3600 lbs.; Rolling Load: 1000 lbs < 0.040". (10 Passes) Impact Load: 150 lbs at 36" drop.

Raised floor Specs (adjacent hallways): Tec-Crete II Screwed Stringer: Concentrated Load: 1500 lbs <0.080"; Uniform Load: 500 < 0.040"; Ultimate Load: 3000 lbs.; Rolling Load: 1500 lbs < 0.040" (10 Passes); Impact Load: 150 lbs at 36".

LARGO
Approximately half of the raised floor area uses cement filled tiles rated at 1,250 psi. If greater loading capacity is needed, the cost of replacing this floor and substructure would be approximately $16K for flooring and $5K labor. This replacement flooring would be rated at (Static Loads) Concentrated Loads = 2500 psi / Uniform loads = 625 psi / Ultimate loads = 6900 psi. *Source http://www.accessfloorsystems.com/ for ConCore 2500 floor tiles.

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 peak power load to other campus buildings was measured at 1070 kW, so Princeton has informed the Government that it may use up to 1430 kW for the Princeton Complex from the current 2500 kW substation. PSE&G has indicated that it will upgrade the capacity of this substation, if required; however, the underground 13200 volt 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 kW, which probably occurred as a result of peak cooling load during the summer probably reflecting attempted start-up of chillers #1 and 3 together or operations during the previous Cray/SGI contract. The peak demand observed on the Main Building substation on March 28, 2005, was 512 kW.

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 actual rating for this transformer is 1725-kVA, but Princeton University mandates that its usage not exceed 1500-kVA. The equipment was installed around 1980 when the Computer Building was constructed and provides power to the Computer Building. Currently it also powers the cooling tower for Chiller #4. Princeton is considering installing an A/B

switch to provide an option to also run Chiller #4 off this transformer in emergency situations. The most recent reading from the Computer Building substation showed a peak load of 654 kW on March 28, 2005.

The Offeror shall be required to install electric meters to measure and record power usage of all of the HPC equipment (i.e., excluding facility infrastructure equipment, such as air handlers, that are operated in the Computer Room), both new and GFE, that are operated under this contract. The Contractor shall be responsible for paying the PSE&G electric utility bill for the Princeton Complex (i.e., power usage associated with both the Main Building and Computer Building transformers). As indicated in Section H.19, the Government will issue contract credits to the Contractor for any power usage above two times the power usage as measured by the above electric meters, where the factor of two includes the power usage associated with cooling the installed system. As indicated in Section C.11.3.2, spreadsheets showing past monthly usage and costs for electric utility bills will be provided to Offerors on request.

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.

All three chillers, described below, currently use the Main Building substation for power.

FAIRMONT

Fairmont has two (2) commercial feeds by Alleghany Power - 4160 V feeders with step down transformers to 480V three phase. All power is fed from a single utility substation with redundant capability. The 480/277-volt main switchboard is rated for 4000 amps. 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 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.

LARGO

The Largo facility receives its power from Potomac Electric and Power Company (PEPCO).

The Largo (ITC) facility has two power feeds. One of these feeds the existing 125kVA UPS and HVAC units and the other fills office/supplementary requirements. Each of these feeds is currently run through a separate Automatic Transfer Switch (ATS), which is connected to a single generator.

The facility manager recommends a separate power feed to the designated rooms, which can be sized as needed to meet the requirements. 150kVA would be available from the generator via separate ATS.

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

Spreadsheets in MS Excel format containing the Princeton Complex's monthly electric utility cost, power usage (in kWH), and peak demand (in kW) for FY1999 through January 2005 will be provided to Offerors upon request.

FAIRMONT

The cost of electrical power has averaged $.05 per KWH over the past several years. How the amount of power consumed is calculated and how it is paid for is negotiable. As stated earlier, power and cooling requirements will have to be accommodated after the requirements are determined.

Total Usage and Expenditures for Electrical Utilities of the FAIRMONT IV&V Facility for FY 2000 - FY 2004

	Annual Usage (kWH)	Expenditure	Avg. Cost / kWH
FY 2000	3,705,600	$184,000	$.0497
FY 2001	3,625,345	$180,000	$.0497
FY 2002	6,554,355	$317,000	$.0484
FY 2003	4,853,702	$235,000	$.0484

LARGO

Based on a recent bill, the total for 14,400 KWH was $2,750.95, which includes Distribution Services, Generation Services, and Transmission Services. This suggests a combined rate of $0.19/KWH.

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 Government/Princeton University), each containing 4 breaker panels, with each panel containing 42 slots
- Liebert 125-KVA PDU (owned by Government/Princeton University), containing 3 breaker panels, with each panel containing 42 slots

- EPE 125-KVA PDU (owned by Government/Princeton University), containing 2 breaker panels, each panel containing 42 slots
  UPS [shown in brown in Figure 3]:
- 500-kVA UPS MGE cabinet and battery bank (owned by Raytheon), Model No. '72-130104-00 EPS 6500/44, 66' with Serial No. 69937-01 (Originally installed in 1997)
- 225-kVA UPS MGE cabinet and battery bank (owned by Government/Princeton University), Model No. '72-130101-01 EPS 6225/44, 66' with Serial No. 200834-01 (Originally installed in 1998)

FAIRMONT
Currently Fairmont has two 1000 kVA UPS stations with associated battery banks to support a 30-minute runtime. Dependence on additional load requirements by the Contractor will determine how many additional PDU's and switch circuits will be available. Any additional requirements above the current capacity will the responsibility of the Contractor.

LARGO
The current UPS and PDU (MGE 75 KVA) are sized to support current operations and some expansion. Therefore, the Contractor must acquire a separate UPS and PDU for the HPC operation based on the probable power requirements.

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 C.11.7.1 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.

FAIRMONT
The NASA IV&V Facility has two diesel-fired, 1250-kW each, generators that activate automatically when power is lost or disrupted.

LARGO
A new Natural Gas Generator is currently being installed. A separate power feed is present with its own Automatic Transfer Switch. This would be connected to the 250KW Natural Gas Generator that is being installed.

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 was completed in 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 permits operation of only one chiller in the Primary Chilled Water Plant at 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 or flow control problems caused by the current chilled water piping design. In particular, Princeton University's experience indicates that Chiller #1 is very sensitive to small changes in flow; as a result, University personnel do not believe that they can operate Chiller #1 along with either of the other two chillers without modifying its flow sensitivity controls. These restrictions could be eliminated by a redesign of the chilled

water plumbing and/or controls; until this situation is resolved, operation appears to be limited to Chillers #3 and #4 for heating loads requiring more than one chiller. Recently, Princeton was able to operate both Chillers #1 and #4 together for a half-hour test; however, attempts to operate Chiller #1 and # 3 continued to fail. The problem with this is that, in the event that Chiller #3 or #4 are unavailable, operation will be limited to only one chiller.

The Government is contracting a cooling engineer to investigate this problem and to determine what needs to be done to provide the chiller redundancy (N+1) that the facility has had in the recent past (using either Chiller #1 or Chillers #3 and 4); this study should be completed by May 2005. It is the Government's intent to correct the problem in FY 2006 using funds outside of the R&D HPCS contract. Since these facility changes will occur after the R&D contract is awarded, the Government intends to coordinate the changes with planned facility modifications proposed by the winning contractor.

Cooling is delivered to the Computer Room through a six-inch and eight-inch piping system from the mechanical rooms 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.

FAIRMONT

The NASA Facility operates three (3) 150-ton chilled-water cooling systems. All of the systems are on generator power. There are three cooling towers, each sized to match the capacity of a single chiller. There is 100% back up (one primary, one redundant) pumping capacity for the chilled water and condenser water systems. Cooling is delivered at 45°F. However, additional cooling capacity may be required to accommodate the system associated with this contract.

LARGO

The ITC has roof-mounted condensers, which feed the two existing HVAC units. There are three condenser/chiller units on the roof of the facility. The facility will be upgrading/replacing HVAC units in March 2005, and additional requirements could be handled in conjunction with those enhancements.

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)
- Seventeen (17) 4-Ton Liebert Extreme Density Overhead (XDO16) Systems
- Cooling powered from UPS

PRTN

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

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

FAIRMONT

The room leased by NOAA contains six (6) Environment Control Units (ECU) that provide the environmental control to the room. However, the current NOAA system occupying the room requires the current capability of all the ECUs. Additional cooling capability will have to be designed and engineered to accommodate the system associated with this contract.

LARGO

The facility has two air handlers. The two air handlers provide N+1 capacity**.** The two HVAC units are Dataguard CCT-DX Series units – total capacity unknown at this time. The units were originally sized to support a UNISYS 1100 Series mainframe and peripheral subsystems. The facility has experience with cooling all systems with a single unit with only one of three compressors operating.

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 (XDO)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.

FAIRMONT

The facility's current chilled water system is running at capacity. Modifications will have to be engineered to accommodate additional requirements.

LARGO

The facility has plumbing that can be used to support water-cooled systems.

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.

Note: If the Contractor chooses to use both BLDR-1 and BLDR-2, any LAN connection between BLDR-1 and BLDR-2 must be made with fiber.

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.

FAIRMONT
Current network connections:

- Verizon OC-3 155 MB link to Camp Springs, MD
- Qwest OC-3 155 MB link to Gaithersburg, MD
- (on order) Qwest Gig-E link to Gaithersburg, MD
- (on order) Qwest Gig-E link to Camp Springs, MD

Note: As these circuits are utilized by NCEP for operational traffic, and are not available for other purposes. However, they are indicative of the telecommunications infrastructure associated with the Fairmont facility.

LARGO
Current network connections:

- Verizon Transparent LAN Service (TLS) -100MB link to the NOAA Silver Spring complex
- Verizon FDDI Network Service (FNS) -10MB link to the Washington area MAN
- 3Mbps WAN (IP over Frame Relay) connected to the four Administrative Support Centers (which are located in Norfolk, Kansas City, Boulder, and Seattle) and to the USDA National Finance Center

This facility is currently served by the earlier FNS 10 MBps and a TLS 100 MBps service. Both are part of the Metro area MAN and connect to Silver Spring and other DC area sites using either FNS or TLS.

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.

FAIRMONT
The facility is located approximately 100 miles from the nearest GigaPop (at the Ohio Supercomputing Center in Pittsburgh).

LARGO
In 2002 the Government investigated the possibility of upgrading the communications resources at the Largo facility. The approximate cost to install a two-pair "dark" fiber link to the University of Maryland Gigapop was estimated by FiberGate to be roughly $80,000 for installation. (FiberGate currently has runs along Landover Road - Route 202, which is 1/2 mile from the ITC.) The monthly recurring maintenance charge does increase, though not in a linear fashion. One pair was estimated at $1500/month and the second pair raised the cost to $2500/month. At this time there are no budgeted funds allocated to implement this upgrade.

Verizon provides network service to this area.

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 Clean Agent fire suppression with VESDA fire detection systems. Both will have wet-pipe overhead sprinklers with 155°F trip point. CO_2 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 coming spring/summer, 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.

FAIRMONT
The computer room is protected by a fire alarm system and a pre-action water sprinkler system.

LARGO
The Largo computer room is protected by a fire alarm system and a wet-pipe water sprinkler system.

C.11.6.2 Fire Extinguishers

BLDR-1 & BLDR-2
Fire extinguishers located within the computer rooms are CO_2 bottles. Fire extinguishers located within the hallways and computer room access areas are FE-36 gas bottles.

PRTN
Fire extinguishers are located at various locations throughout the Computer Room and in the Storage Room adjacent to the Loading Dock. 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.

FAIRMONT
Fire extinguishers are located within the computer rooms and throughout the facility. The fire extinguishers located within the computer room are designed for use with electronic equipment.

LARGO

Fire extinguishers are located within the computer rooms and throughout the facility. The fire extinguishers located within the computer room are designed for use with electronic equipment.

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 (Plainsboro Fire Department as first responders, backed up by Princeton Plasma Physics Lab) 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.]

FAIRMONT
THE NASA IV&V Facility is supported by the Fairmont Fire Department. The average response time is approximately 7 minutes. Each year the members of the fire department are given a tour of the facility to keep them familiar with the facility and to keep them informed of what is within the facility.

LARGO
In the most recent incident the local fire department arrived on scene within 10-15 minutes of being notified. The closest fire department, Kentland Volunteer Fire Department, is located less than five miles from the facility.

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**	**Resulting Downtime**
26 Mar 03	13 Minutes	0 Min
13 May 03	13 Minutes	0 Min
8 Sep 03	30 Seconds	0 Min
11 Sep 03	2 Minutes	0 Min
FY2004		
10 Jun 04	52 Minutes	0 Min

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, a review of operator logs showed 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

FAIRMONT
All critical electrical systems are fully backed up for uninterrupted service to cover for an extended length of power outage. The facility's operational systems are protected by UPS systems that are supported by a bank of batteries and two diesel generators. Although the diesel generators come on line automatically within one minute if a power loss or disruption is detected, the battery system will provide uninterrupted power for up to 30 minutes.

LARGO
A new natural gas generator is currently being installed. There have been five power outages in CY 2004. Three of these were less than one hour in duration.

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
Based on panel read-outs, the UPS systems for the current system, which is in transition to the April 2005 upgrade, provide a ride-through for the computer systems of 16 minutes on the 500-KVA UPS and 25 minutes on the 225-KVA UPS. The 16-minute duration for the 500-kVA UPS is due to one of the four wet-cell battery banks being over 4 years old; this bank is expected to be upgraded by the first week of April 2005, resulting in all of the batteries in this UPS being less than 6-months old. This upgrade should produce a ride-through time close to 25 minutes. The wet-cell batteries in the 225-kVA UPS are all less than 6 months old. Once the April 2005 HPCS installation is complete, the UPS capacities (at 386 and 109 kVA for the 500- and 225-kVA UPSs respectively on March 23) are expected to both be at maximum (80%) capacity, and the ride-through times are likely to be reduced somewhat. (Note: The PDUs are also expected to be at full capacity after the April 2005 installation is complete.) The room air conditioning units and the chillers do not have generator or UPS 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.

FAIRMONT
The NASA IV&V Facility is designed to provide uninterrupted power to all current critical operational systems.

LARGO
There is no UPS associated with the proposed facilities.

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.

FAIRMONT
The Fairmont center currently has remote dial-up capability, and is expecting an upgrade to have full remote monitoring capability within the next six months.

LARGO
The computer room does not have remote monitoring capabilities.

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 94F.

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

FAIRMONT
The NASA IV&V Facility is typically a Monday – Friday, 7:00 AM – 6:00 PM operation, but Security personnel are present 24/7/365; therefore, after-hours access is available.

LARGO
The facility at Largo is manned round-the-clock on every day of the year.

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 patrol the Forrestal Campus on a regular basis.

FAIRMONT
NASA security personnel are located on site 24/7/365.

LARGO
Federal Protective Service patrols on a regular basis.

C.11.8.2 Computer Room Security

BLDR-1 & BLDR-2
Computer room has restricted access provided by proximity keypads and other controls.

PRTN
Computer Room access is controlled by a cipher lock system and other controls.

FAIRMONT
The computer room has restricted access provided by proximity badge readers and other controls. Because the facility is used by a government organization, access to the facility is controlled by NASA per NASA regulations.

LARGO
The computer room has restricted access provided by proximity badge readers and other controls.

C.11.8.3 Procedures for Contract Personnel to Access a Government Facility

BLDR-1 & BLDR-2 & PRTN & FAIRMONT & LARGO
All contract personnel must submit to a background check and will be issued an "affiliate" ID badge. Authorized NOAA and contractor personnel may escort other NOAA and contractor personnel as long as the visitors are US citizens. Special arrangements must be made weeks in advance to accommodate escorted access by foreign nationals.

FAIRMONT
Access to the NOAA room is authorized by NOAA but controlled by NASA IV&V Facility Security personnel and systems. Authorized NOAA and contractor personnel may escort other NOAA and contractor personnel into the NOAA area as long as they are US citizens. Special arrangements must be made weeks in advance to accommodate escorted access by foreign nationals.

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.

FAIRMONT
Annual facility costs are approximately $54 per sq. ft. per year, which includes the following: most utility costs (but not the mass consumption of electricity); some ADP support; security; and facility support, such as cleaning, mail delivery, etc. Mass electrical usage is calculated and charged separately at approximately $.05 per kWH.

LARGO
Annual facility costs, excluding electrical utility costs, are estimated to be $28/sq. ft. This covers the entire facility cost, including water and electrical power, with the addition of a charge of approximate $45K/year for after normal business hours and weekend power.. The ITC facility is 16,250 sq. ft., of which 15,950 are usable (water heater, etc. makes up the difference).

C.11.10 Availability of Blueprints

BLDR-1 & BLDR-2 & PRTN & FAIRMONT & LARGO
Blueprints of the facility are available and will be provided to Offerors, subject to their approval of conditions for their use of the drawings.

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.

		Oct. 2004	Oct. 2005	Oct. 2006
Available Raised Floor Space (sq. ft.)				
	BLDR	0	1,424	3,674
	PRTN	900	2,700	5,500
	FAIRMONT	3,200	3,200	3,200
	LARGO	920	920	920
	TOTAL	***6,020***	***8,244***	***13,294***
Available Non-Raised Floor Space (sq.ft.)				
	BLDR	0	0	0
	PRTN	360	900	2,700
	FAIRMONT	0	0	0
	LARGO	724	864/1,440	864/1,440
	TOTAL	***1,084***	***1,224/1,800***	***1,224/1,800***
Available Power for New Hardware (kVA)				
	BLDR	0	350	600
	PRTN	250	150	575
	FAIRMONT	150	150	150
	LARGO	***0***	150	150
	TOTAL	***400***	***800***	***1,475***

		Oct. 2004	Oct. 2005	Oct. 2006
Available Cooling for New Hardware (tons)				
	BLDR	0	146	218
	PRTN	180	155	283
	FAIRMONT	0	0	0
	LARGO	0	*	*
	TOTAL	***6,180***	***8,480***	***13,530***
*NOTE: The available cooling specs for the Largo facility will be provided to Offerors during the scheduled site visit.				

Assumptions for Projections

Available Raised Floor Space:

BLDR:
(1) October 2005 estimate of 1424 ft. is solely in BLDR-2 and assumes completion of computer room design, construction and testing.
(2) The increase of 2250 sq. ft. in the October 2006 total reflects additional available space in BLDR-1 that is currently 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 reflects that the StorageTek silos are still located on the raised floor.
(2) Oct. 2005 number is an estimate derived from the unoccupied raised floor space for the room layout as indicated in Figure 3a, which shows the assumed positions of systems and Storage Tek silos after the mid-contract installation is complete. (The actual square footage of free space is obviously subject to interpretation to account for obstructions such as room columns and ramps, size of equipment to be installed, etc.)
(3) Oct. 2006 estimate reflects available space occupied by systems leased under current Raytheon contract, because this leased equipment will be removed at the end of the current 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.

FAIRMONT:
The available space in the Fairmont facility is available for use in July 2006.

LARGO:
The available space in the Largo facility is available for use in March 2005.

Available Non-Raised Floor Space:
PRTN:
(1.) Oct. 2005 estimate assumes a common staging area of 360 sq. ft. on the non-raised floor adjacent to the exit to the loading dock plus an area of 18'x30' on the left hard pan that would be available to the R&D Contractor for office space and storage. The remaining 18'x32' of non-raised floor space is reserved for storage and work area until the end of the current Raytheon contract.
(2.) Oct. 2006 estimate assumes both the left and right hard pan areas of the Computer Room are available to the R&D Contractor, with the square footage shown including the UPS Room as well. As with the raised floor space, the timing for availability of non-raised floor space requires coordination with the current contractor to assure smooth transition from the previous contract to the new R&D contract.

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.

FAIRMONT:
It is the assumption that NOAA and NASA will have to make modifications to the facility to meet the power requirements and could be available within 6 months of the requirements being established. Studies are being performed now to determine alternate approaches.

LARGO:
The Government requires that the Contractor acquire a separate UPS and PDU to support any HPC components that are placed in the LARGO facility.

Available Cooling:

BLDR:
(1.) The chilled water plant has reached its maximum capacity and will lose its N+1 capability on hot days. A failure within the chilled 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 a factor of 1.2 to reflect (conservatively) other room heat sources and cooling inefficiencies.
(3.) Assume maximum chiller capacity uses chillers #3 (350 tons) and #4 (225 tons) operating together. Joint operation of chillers #1 (400 tons) and #4 (225 tons) are not currently viable as explained in Section C.11.4. However, as indicated in Section C.11.4, joint operation of chillers #1 and #3 is expected to require mechanical (plumbing) modifications.

FAIRMONT:
The chilled water plant has reached its maximum capacity. It is the assumption that NOAA and NASA will have to make modifications to the facility to meet the cooling requirements and could be available within one year of the requirements being established. Studies are being performed now to determine alternate approaches.

LARGO:
The assumption is that the facility has adequate cooling available. However, this is dependent on the size and type of equipment that the contractor may chose to place at this facility.

		Oct. 2004	Oct. 2005	Oct. 2006
Available Wide-Area Network Capacity (Mb/sec.) (list by type below)				
	BLDR			
Abilene / Internet 2		310	500	0
Commercial Internet		80	80	100
National Lambda Rail (NLR)				
NOAA Research / Internet 2		0	1,000	10,000
Commercial Internet		0	1,000	1,000
***Note: 10,000 Mbps and/or dedicated Lambda's (wavelengths) may be available via NLR in 2005 And 2006 if needed.**				
	PRTN			
Microwave to Internet2 (PU allowed)= 50 Mb/s		25	20	15
Commercial Internet = 9 Mb/s		3	3	3
	FAIRMONT			
No Government-furnished WAN capacity available		0	0	0
	LARGO			
Verizon Transparent LAN Service		100	100	100
	TOTAL	***515***	***2,700***	***11,215***

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.

FAIRMONT:
There is no government-furnished WAN capacity available for the contract. Therefore, the Offeror will need to provide any WAN capacity required for the contract.

LARGO:
The WAN capacity might need to be upgraded depending on the HPC components that are placed in this facility.

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

Figure 3a *Computer Room Layout for PRTN as of April 2005. Equipment and room features (columns, etc.) that are expected to remain on the raised floor after October 2006 are color-coded in the figure, although the Contractor may choose to reposition or remove air-handlers, PDUs, etc.*

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

Figure 4 *Computer Building Schematic for Princeton Complex*

[Figure 5 withheld from Public Website – Available upon request.]

Figure 5 *Schematic of Computer Rooms for LARGO*

[Figure 6 withheld from Public Website – Available upon request.]

Figure 6 *Schematic of Computer Room for FAIRMONT*

Request for these figures are to be submitted to the Contracting Officer at the following e-mail address: william.voitk@noaa.gov

C.12 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
Chloride	EDP90/300/250 4x4	Uninterruptible Power Supply UPS Cabinet and Battery Bank (250kVA)	1	10/2006
	EDP70L/100/100 4x4	UPS Cabinet and Battery Bank (100kVA)	2	10/2006
Liebert	FH740C	Computer Room Air Conditioning Chill Water - Down Draft – 30 tons	4	10/2006
Phonetics	SCADA 3000	Environmental Monitoring Sensaphone SCADA 3000	1	10/2006

Unrestricted GFE

Manufacturer	Part/ Model Number	Description	Qty	Date Available
		Compute Hardware		
Myricom	M3-E128	Myrinet line card enclosure	15	10/2006
Myricom	M3-M	Management line card	15	10/2006
Myricom	M3-SW16-4DM	Ribbon spine line card	84	10/2006
Myricom	M3-SPINE-8F	Fiber spine line card	44	10/2006
Myricom	M3-SW16-8F	Fiber switch line card	96	10/2006
Myricom	M3F-PCI64B-2	PCI card (NIC)	768	10/2006
Myricom	M3-CLOS-ENCL	14U enclosure for Clos256_256 networks	2	10/2006
Myricom	M3F-PCIXD-2	PCI D card (NIC)	317	10/2006
Myricom	M3-SW32-16F	16 Port line card	20	10/2006
Myricom	M3-THRU-16Q	Line card for 16 quad fiber thru connections	2	10/2006
Myricom	M3-2SW32	Dual SW32 line card	4	10/2006
Myricom	M3-4SW32-16Q	Quad SW32 line card	2	10/2006
		HSMS		
ADIC	AML/J	Tape Robot	1	10/2006
IBM		LTO-1 Tape drives	8	10/2006

Manufacturer	Part/ Model Number	Description	Qty	Date Available
Sun	E450	HSMS Server	1	10/2006

Washington, DC

All of the GFE listed in this subsection below will be available October 2006.

Site-Constrained GFE

- None

Unrestricted GFE

Gaithersburg, MD:
Two (2) StorageTek Powderhorn silos configuration (two 9310, 9311, 9330).
Each of the silos includes:

- eight (8) 9940B tape transports
- 5000 9940 tape cartridges

Each silo contains roughly one PetaByte of data.

Fairmont, WV:
One (1) StorageTek Powderhorn silo configuration (9310, 9311, 9330).
The silo includes:

- four (4) 9940B tape transports
- 2300 tape 9940 cartridges

Princeton, NJ

Unless otherwise stated, all of the following GFE for Princeton will be available in October 2006

Site-Constrained GFE

Manufacturer	Part Number	Description	Qty
		Computer Room Security Monitoring	
		C-Cure 800	
Compac	4403US	Compac MS 2000 C-Cure 800 OS	1
Emerson	P761VBT0EENC	Monitor C-Cure 800	1
		Computer Room Security Monitoring	
		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

		Power Distribution Units	
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
		Uninterruptible Power Supply	
MGE	72-131522-000	UPS Cabinet and Battery Bank 225 KVA	1
		Power Distribution Units	
Liebert	PPA125C	PDU Unit (125 KVA)	1
EPE	PD084M48A12-125	PDU Unit (125 KVA)	1
		Computer Room Air Conditioners	
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

Manufacturer	Part Number	Description	Qty
		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 rsponse, 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)	16
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

C.13 Appendix D – Terms

C.13.1 Abbreviations

ADIC	Advanced Digital Information Corporation
ANSI	American National Standards Institute
CM	Configuration Management
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 (10^9 bytes)
dPB	Decimal Petabyte (10^15 bytes)
DOC	US Department of Commerce http://www.commerce.gov
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 http://www.fsl.noaa.gov 325 Broadway Boulder, CO 80305
FY20xx	Fiscal Year 20xx
Gb	Gigabit (2^30 bits)
Gb/s	Gigabit(s) per second
GB	Gigabyte (2^30 bytes)
GFDL	NOAA's Geophysical Fluid Dynamics Laboratory http://www.gfdl.noaa.gov 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

HIS	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 (2^10 bytes)
KVA	Kilovolt amperes (1000-volt amps).
LAN	Local Area Network
LPAR	Logical Partitioning
LSC	Large-Scale Computing
LTD	Live Test Demonstration
MAN	Metropolitan Area Network
Mb	Megabit (2^20 bits)
Mb/s	Megabit(s) per second
MB	Megabyte (2^20 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 http://www.ncep.noaa.gov 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 http://www.noaa.gov
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 (2^50 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
SLT	System Life Throughput
SON	Statement of Need (Section C of this Procurement)
STK	Storage Technology
TB	Terabyte (2^40 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
WS	Workstream

C.13.2 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. See Section C.6.1.1.
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.
Job Slot	The logical partition of a computational resource supporting an instance of a workstream (e.g. the set of processing elements on which an instance of a workstream is run).

Null Time	The period of time when the workload cannot be accomplished due to environmental failure at a Government furnished site, 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	The set of components which accomplish a workstream suite. The aggregate of all Subsystems is the HPCS. Synonym for Target IT Architecture.
System Life Throughput (SLT)	A measure of performance and availability delivered for all instances of a given workstream. See Section C.6.1.2 for more information.
Target IT Architecture	The set of hardware and software components 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.