BAA 06-29
Proposer Information Pamphlet (PIP)

for

Defense Advanced Research Projects Agency (DARPA)
Strategic Technology Office (STO)

Dynamic Multi-Terabit Core Optical Networks: Architecture, Protocols, Control and Management (CORONET)

Technical POC: Dr. Adel Saleh, DARPA/STO

****
This BAA will be open for **three (3) years** from the date of its publication in www.fbo.gov.

NOTE: Although this BAA will be open for three (3) years from the date of its publication on www.fbo.gov, the Government anticipates that the funding for this program will be committed only during the first and second selection periods. To be considered for funding during the first round of selections, for Phase 1, full proposals must be received no later than **12:00 NOON Eastern Daylight Saving Time on October 27, 2006.** The proposal window for the second round of selections, for Phase 2, will be announced during the latter half of the Phase 1 execution.

A Proposers' Day Conference will be held on August 8, 2006 to encourage discussion and teaming on this topic. Additional information is contained inside this PIP. Questions and Answers will be subsequently posted to http://www.darpa.mil/ato/solicit/CORONET/index.htm for Proposer review. NOTE: It is not mandatory to attend the Proposers' Day Conference to respond to this BAA.
****

BAA 06-29 CORONET

1. INTRODUCTION

The Defense Advanced Research Projects Agency's (DARPA) Strategic Technology Office (STO) is soliciting proposals under this BAA for the performance of research, development, design, and testing to support the DARPA Dynamic Multi-Terabit Core Optical Networks: Architecture, Protocols, Control and Management (CORONET) program.

1.1. APPROACH

This BAA affords proposers the choice of submitting proposals for the award of a Procurement Contract, Technology Investment Agreement, Other Transaction for Prototype Agreement, or other such appropriate award instrument. The Government reserves the right to negotiate the type of award instrument determined appropriate under the circumstances.

1.2. PROPOSERS

The Government encourages proposals from non-traditional defense contractors, nonprofit organizations, educational institutions, small businesses, small disadvantaged business concerns, Historically-Black Colleges and Universities (HBCU), Minority Institutions (MI), large businesses and Government laboratories. Teaming arrangements between and among these groups are encouraged. However, no portion of this BAA will be set aside for HBCU/MI, small, or small disadvantaged business participation due to the impracticality of preserving discrete or severable areas of research in the technologies sought. Government/National laboratory proposals may be subject to applicable direct competition limitations, though certain Federally Funded Research and Development Centers are excepted per P.L. 103-337 § 217 and P.L. 105-261 § 3136. Any responsible and otherwise qualified Proposer is encouraged to respond.

Proposers for Phase 1 and Phase 2 must be U.S. companies. Contractor teams for Phase 1 may include foreign firms, foreign universities and foreign individuals so long as: the prime is a US company, team member firms or universities meet the criteria of this BAA, the U.S. Government is otherwise permitted to conduct business with the foreign firm or university and all foreign personnel sign all appropriate non-disclosure agreements prior to participating in the research effort. It is possible that the Phase 2 effort, or portions thereof, will significantly limit foreign participation in accordance with U.S. Export Control Laws.

1.3. PROGRAM SCOPE AND FUNDING

The Government intends to issue awards based on the optimum combination of proposals that offers the best overall value to the Government. The Government reserves the right to award without discussions. The Government reserves the right to select for award all, some, or none of the proposals received in response to this BAA. The Government also reserves the right to select for award some portion(s) of the proposals received; in that event, the Government may select for negotiation all, or portions, of a given proposal. The Government may incrementally fund any award issued under this BAA.

It is anticipated that this effort will be conducted in two phases. Proposers should propose a complete and comprehensive solution for Phase 1 and separate proposal for Phase 2 and submit them during the appropriate selection periods. Proposals addressing only a subset of the topics will be considered non-responsive to this BAA. As such, teaming of proposers to assemble the necessary breadth and depth of expertise and resources is recommended. To be considered for funding during first round of selections, for Phase 1, full proposals must be received no later than 12:00 NOON Eastern Daylight Saving Time on October 27, 2006. The proposal window for the second round of selections, for Phase 2, will be announced during the latter half of the Phase 1 execution.

While the earliest anticipated award is planned to occur in December 2006, the Government may select for funding any full proposal or portions of a proposal at any time while this BAA is open. It is expected, however, that no funding will be available for any proposals submitted before, between, or after the first and second selection periods; and therefore, the Government does not anticipate awarding any contracts outside the two selection windows.

1.4. PERIOD OF PERFORMANCE

The total period of performance for the effort is incremental and is planned as follows:

- Phase 1: Effective date of award through eighteen months (Months 1-18)
- Phase 2: If awarded, will be from the date of award through completion of validation testing; the total period for this phase has not been determined.

The Government may incrementally fund any awards under this BAA.

1.5. TECHNICAL AND ADMINISTRATIVE SUPPORT

It is the intent of this office to use contractor support personnel in the administration of all submittals to this BAA. The Government intends to use non-government employees and subcontractors, to include, but not limited to SRS Technologies, Arlington, VA, SAGE Solutions Group Inc., Centreville, VA, MIT – Lincoln Labs, Cambridge, MA, and Booz Allen Hamilton, Arlington, VA to assist in administration and, if needed, provide technical expertise on portions of the proposals. These personnel will have signed and be subject to the terms and conditions of non-disclosure agreements. By submission of its proposal, a proposer agrees that its proposal information may be disclosed to employees of these organizations for the limited purpose stated above. If you do not send notice of objection to this arrangement, the Government will assume you consent to use the subject personnel in review of your submittal(s) under this BAA. Only Government personnel will make technical evaluations and award recommendations or decisions under this BAA.

1.6. INSTRUCTIONS AND POINTS OF CONTACT

All questions pertaining to this BAA may be submitted to DARPA at this e-mail address: **BAA06-29**@darpa.mil. DARPA may post updates to questions or comments periodically to the solicitation website: http://www.darpa.mil/ato/solicit/CORONET/index.htm

2. OVERVIEW OF THE CORONET PROGRAM

2.1. PROGRAM SUMMARY

The objective of the CORONET program is to revolutionize the operation, performance, survivability and security of the United States' global inter-networking infrastructure. In particular, proposals are sought that advance the network architecture, protocols, and control and management for highly dynamic, multi-terabit core optical networks.

The program targets a global network that should enable up to a ten-fold increase in aggregate network demand over today's state-of-the-art networks while sharply reducing the cost, size, and power requirements per bit, and greatly simplifying network operations. The network should support very fast service set-up/tear-down to allow low-latency, collaborative and grid-computing applications across the globe. The network also should enable very fast and efficient recovery from multiple concurrent network failures.

The focus of this program is the core network, which is the large-capacity long-haul infrastructure that interconnects the major sites of the global network. However, optimizations of the architecture and protocols of the edge networks, which connect users to the core, can be considered in order to improve the overall network performance and economy. Proposals that consider an integrated edge/core network design approach are encouraged.

Delivering a cost-effective, capacity-efficient solution that can support an order-of-magnitude increase in network demand and a high degree of network agility and resiliency will require new paradigms in areas such as cross-layer traffic processing and capacity management, as well as significant advancements in network operation. A major thrust of the program is to develop the architecture, protocols, algorithms, and control and management systems that will enable the desired advanced network capabilities. The performance of these developed systems will be evaluated through extensive network modeling, simulations, and demonstrations.

It is envisioned that the majority of traffic carried on the network will be Internet Protocol (IP) carried directly over optical wavelength division multiplexing (WDM). However, a substantial amount of traffic will be wavelength services (i.e., end-to-end circuits consuming an entire wavelength or a multiplicity of wavelengths). For purposes of evaluating proposed solutions, a specific set of network services is specified in Section 2.2. However, it is important that proposals be robust to the exact mix and underlying composition of the network traffic.

The details of the target network topology and assumptions about the underlying physical layer properties are specified in Section 2.2. Any proposed solution must be scalable; solutions should be feasible for near-term operation, as well as operation in an environment with ten times the amount of traffic as today's networks. To probe the scalability of proposed solutions, performers must model several network growth scenarios, as specified in Section 2.2, where the offered load (and the underlying capabilities of the network) steadily increases in each setting.

Assumptions regarding network technology are also provided in Section 2.2. It is assumed that the network utilizes all-optical switching elements and ultra-long-reach all-optical transport. This assumption is made because such technologies have been proven commercially to be more scalable in cost, size and power requirements than traditional optical-electrical-optical (OEO) based solutions in large core networks. However, solutions based on other technologies also can be considered as long as they are scalable in cost, size and power, for both transport and switching.

In addition to serving as the next-generation global inter-networking infrastructure for the Department of Defense, it is essential that the results of this program be suitable for commercial deployment as well. The target deployment time frame is about five years in the future from contract award date.

In order for the Government to evaluate the effectiveness of proposed solutions in achieving the stated program objectives, a set of performance metrics is specified in Section 2.5. These metrics must be met for each of the design scenarios provided in Section 2.2. In addition, performers are encouraged to identify other characteristics relevant to their particular network design that they propose to quantitatively evaluate.

A network study exercise is included as part of the proposal, to assess the capabilities and networking expertise of proposers and gauge preliminary algorithmic performance with respect to network resiliency.

In Phase 1 of the project, performers must design the architecture for the data, control and management planes, and the necessary protocols and algorithms to enable the desired service requirements. Performers must develop appropriate-level-fidelity models of their designs, and must perform network studies for each of the design scenarios in Section 2.2. Performers must deliver valid and verifiable simulation results that meet the specified program metrics. The period of performance of Phase 1 is 18 months.

As described in Section 2.4.2, one of the deliverables at the end of Phase 1 is a full (technical and cost) proposal for Phase 2 detailing a proposed development program for a network control and management software system implementing the architecture, protocols and algorithms devised in Phase 1. The proposal should include milestones, metrics, and a comprehensive test plan. An important objective of Phase 2 is to demonstrate the advanced services and enhanced operational features enabled by the developed system, such that the results of the CORONET program are attractive for transition to commercial carriers.

A high-level proposal and rough order of magnitude (ROM) cost for Phase 2 should be included in the proposal responding to Phase 1 of this BAA outlining the proposed Phase 2 development program, test plan and strategy for commercial transition.

2.2. PROBLEM DEFINITION

Section 2.2.1 specifies the key network parameters and attributes for the CORONET program. These are the parameters to be assumed for each of the required network studies. Section 2.2.2 provides a high-level description of two of the major technological challenges of the program: network agility and network resiliency.

2.2.1. Network Attributes and Parameters

2.2.1.1. Topology Attributes

The target network is a globally distributed core optical network, encompassing nodes in both the Continental United States (CONUS) and various sites across the globe. Nodes are the sites in the network where terminals and switches are located, where traffic is added, dropped, switched, and/or groomed. The total number of network nodes, as well as the distribution of nodes inside/outside CONUS, is specified below, in Table 1.

Network links are the physical fiber connections between adjacent nodes. The degree of a node refers to the number of network links (i.e., line-side links) incident on a node (with each link populated by one or more fiber pairs). In realistic core networks of size comparable to the target network, the majority of nodes have a degree of two, with very few nodes having a degree greater than four. For purposes of this program, the percentage of nodes having a particular degree is fixed in Table 1. (Any undersea links incident on a node should be counted when determining the nodal degree.) Furthermore, the rule that is to be used for specifying the link distance between adjacent nodes is provided in Table 1.

It is crucial that a high degree of survivability be achievable for some subset of the traffic. Thus, it is important that there be multiple link/node diverse paths between traffic endpoints. The assumptions for topology diversity are stated in Table 1.

Table 1. General Topological Attributes of the Target Optical Network

Parameter or Item	Assumptions
Number of Nodes in the Core Optical Network	100 nodes in total, of which 75% are within CONUS and 25% are distributed globally.
Nodal Degree Distribution	50% of the nodes are degree-2 nodes, 35% degree-3 nodes and 15% degree-4 nodes. If needed, up to 3% of the nodes can be changed to degree-5 or degree-6 nodes.
Distance between Adjacent Nodes	The fiber link distance between adjacent nodes is assumed to be 20% longer than the direct 'as-the-crow flies' distance between the nodes.
Fiber Topology Diversity	Three diverse east-west paths across CONUS. CONUS is assumed to be connected to the global network via stand-alone undersea systems with three diverse cables in each of the Atlantic and Pacific oceans, with sufficient capacities. Other undersea cables across the globe can be assumed as needed.
Node Sizes and Traffic Distribution	A realistic, i.e., non-uniform, traffic should be assumed. As a guide, if 20% of the nodes are assumed to be "large" nodes, and the rest "small" nodes, the traffic distribution among the two types of nodes should be approximately: Large-Large: 40%, Large-Small: 40%, Small-Small: 20%

It is incumbent upon the performers to generate an appropriate global network topology that meets these requirements. The resultant topology is to be used in the network design exercise for the proposal (see Section 4.2.1.1). An illustration of the network must be included as part of the proposal, and the realism of the network topology will be an important factor in evaluating the proposal. This network topology also will be used in the network design studies of Phase 1, although small changes to the topology will be permitted, if necessary.

It is also important that performers demonstrate realism in generating a traffic distribution to meet the specifications of the CORONET program. All nodes in a core network do not source an equal amount of traffic. Typically, there are a small number of nodes that generate a large percentage of the traffic. Guidelines for nodal size are provided in Table 1. Traffic distributions based on 'all-to-all' traffic are not admissible. A sample traffic distribution used for the initial network study must be included as part of the proposal (see Section 4.2.1.1).

While performers should assume the attributes stipulated in Table 1 for purposes of evaluation, it is important that any proposed solution be robust with respect to perturbations from these specified characteristics.

2.2.1.2. <u>Network Services</u>

The level of IP traffic in current networks is steadily increasing, and in many networks already represents a majority of the network traffic. IP is used for a range of services, such as web browsing, file transfers, telephony, streaming video, interactive video, etc. As a result, the requirements for IP flows are highly divergent. In this section, we focus on the data-rate characteristics of the IP traffic. Metrics for quality-of-service (QoS) attributes for IP traffic are specified in Section 2.5.3.

Some IP services, such as Voice Over IP (VOIP), IPTV, and other services based on circuit emulation over IP, require guaranteed bandwidth from the network in order to achieve an acceptable level of performance. Such guaranteed-bandwidth services can be roughly categorized as fine granularity, medium granularity, and coarse granularity, depending on the amount of required bandwidth. For purposes of this program, it can be assumed that the data rates of guaranteed-bandwidth services have a very small variance; i.e., the data rate of such services remains close to the contracted amount.

In contrast to guaranteed-bandwidth services, best-effort flows are not provided with individual minimum guaranteed bandwidth. This class of service is used for non-time-critical applications such as e-mail and web browsing. The data rate characteristic of such flows tends to be fairly bursty.

In addition to a predominance of IP traffic, it is expected that there will be a substantial amount of wavelength-service traffic as well. In such services, a circuit consuming an entire wavelength, or a multiplicity of wavelengths, is requested for the duration of the connection.

The percentage of IP services and wavelength services to be assumed in the target CORONET network is shown in Table 2, as a function of the parameter S. This parameter represents the percentage of the aggregate network demand that is IP services, with the remainder being wavelength services. In the network studies that are to be performed as part of the CORONET program, S ranges from 50% to 100%. The nominal value of S, denoted by S_o, is 75% (this is the value to be assumed whenever the PIP refers to a nominal network design).

Table 2. Network Services

High-Level Service Type	IP Services				Wavelength Services [3]			
High-Level Traffic Breakdown [1]	$S \in [50\%, 100\%]$ The nominal value $S_o = 75\%$				1-S The nominal value $1\text{-}S_o = 25\%$			
Service Sub-Type	Best-Effort IP Services	Guaranteed-Bandwidth IP Services			Single-λ Services	Double-λ Services	Quad-λ Services	Octal-λ Services
		Fine Granularity	Medium Granularity	Coarse Granularity				
Service Data Rate [2]	Bursty	< 400 Mb/s	0.4 to 4 Gb/s	4 Gb/s to 0.25 R	R	2 R	4 R	8 R
Service Sub-Type Traffic Breakdown	40% of S	15% of S	20% of S	25% of S	40% of (1-S)	20% of (1-S)	20% of (1-S)	20% of (1-S)

(1) This row represents the fraction of the aggregate network demand that is in each high-level service category, where the aggregate network demand is provided in Table 3 for four design scenarios. This row is specified in terms of *S*, where *S* can take on a range of values, as specified.

(2) The data rates of some of the service types are specified in terms of *R*, the bit rate per wavelength, which is provided in Table 3 for four design scenarios.

(3) Single- or multi-wavelength services are digital-circuit services of the data rate specified.

Both IP and wavelength services are further broken down by sub-type in Table 2. The fraction of aggregate demand that is in each traffic sub-type is again a function of *S*. IP services are divided into best-effort services, as well as fine, medium, and coarse granularity guaranteed-bandwidth services. Note that the data rate of an individual guaranteed-bandwidth IP flow is less (typically much less) than 25% of a wavelength. Efficiently carrying such services in the optical layer is an important area of research in the CORONET program. Wavelength services are broken into single-, double-, quad-, and octal-wavelength services, which require one, two, four and eight wavelengths' worth of bandwidth, respectively.

The data rate of some of the service sub-types is specified in terms of the parameter *R*, which represents the bit rate per wavelength (i.e., the 'line rate') for a particular design scenario, as will be described in Section 2.2.1.3.

Also, note that all percentages in this table (as well as in the remaining tables) are based on the *total bandwidth* of the aggregate network demand, and not on the *total number* of demands. The aggregate network demand for each design scenario will be specified in Section 2.2.1.3.

For purposes of evaluation against the metrics of Section 2.5, all demands should be assumed point-to-point and bi-directionally symmetric. However, support for asymmetric services and multicast services must be included as part of the proposed solution.

Note that Table 2 implicitly includes value-added network services such as Private Line. It is assumed that, for example, Private Line services are mapped into guaranteed-bandwidth IP services. While the CORONET program specifies an IP-over-WDM architecture, it is recognized that this is not the only model possible for carrying future services. In addition to analyzing the IP-over-WDM model, as specified in this PIP, performers are permitted to also propose and analyze other models if they can provide convincing data that these models offer significantly higher performance with comparable or lower cost, or comparable or higher performance with significantly lower cost.

2.2.1.3. Design Scenarios

Current state-of-the-art core networks are designed for an aggregate network demand on the order of 10 Tb/s, where aggregate network demand is defined as the sum of all bi-directional end-to-end demands on the network. As data-intensive applications such as grid computing, remote visualization, and real-time fusion and synchronization of sensor data continue to grow, it is important that the next-generation core network be able to handle up to an order-of-magnitude increase in demand.

This ten-fold growth is expected to emerge over the next 5 to 10 years. Any proposed solution must be practicable as the network grows over this time frame. In order to test the scalability of proposed solutions, a set of evolutionary demand scenarios is specified in Table 3. Each scenario specifies an aggregate network demand, as well as the corresponding physical layer assumptions. Given that the target network will support dynamic traffic, with some connections being established and torn down in a short time frame, and given the bursty nature of best-effort IP services, the aggregate network demand specified in Table 3 should be interpreted as the long-term average demand for the given scenario.

Note that the aggregate network demand represents the total demand in the global network. Some of this demand will be contained solely within CONUS, some will be strictly outside of CONUS, and some demands will cross the CONUS boundary. Reasonable assumptions should be made to apportion the demands across these categories.

The network attributes and services specified in Table 1 and Table 2 hold for each of the design scenarios of Table 3. The parameter *R*, referenced in Table 2, is specified for each scenario in the row 'Bit Rate Per Wavelength'. For each scenario, performers must meet the metrics that are specified in Section 2.5. In addition, it is expected that proposed solutions will exhibit a decrease in cost, size, and power requirements per bit as the network grows.

The capacity per network link is assumed to increase as the aggregate demand increases. This is accomplished by increasing either the number of fiber pairs per network link or the bit rate per wavelength, as indicated in Table 3. Increasing the bit rate per wavelength, while keeping the number of wavelengths per fiber constant, represents an increase in the system spectral efficiency. Performers should assume that such technological advances would be available as the network evolves. Research regarding schemes to increase the system spectral efficiency is outside the scope of this program.

Table 3 specifies a single line rate for each demand scenario. However, any proposed solution should be compatible with a hybrid line-rate architecture (e.g., a combination of 40 Gb/s and 100 Gb/s wavelengths), which may occur as the network evolves.

Table 3. Scenarios for Aggregate Network Demand, Bit Rate per Wavelength, and Fiber and Network-Link Capacities

	Scenario 1	**Scenario 2**	**Scenario 3**	**Scenario 4**
Aggregate Network Demand [(1)] (**Tb/s**)	20	40	50	100
Bit Rate per Wavelength (**R**) (**Gb/s**)	40	40	100	100
Capacity per Fiber (**Tb/s**)	4	4	10	10
Number of Wavelengths Per Fiber	100	100	100	100
Nominal Number of Fiber Pairs per Network Link [(2)]	1	2	1	2

(1) The *Aggregate Network Demand* is defined as the sum of all bi-directional end-to-end demands on the network.

(2) Ideally, the aggregate network demand is carried using no more than the number of fiber pairs per link specified for each scenario. However, if necessary, one extra fiber pair can be added on heavily loaded network links as specified in Table 6. The undersea cables associated with each scenario are assumed to have sufficient capacity.

2.2.1.4. Optical Hardware Technology

In legacy core networks, the optical signal carried by each wavelength is converted to the electronic domain at every network node along the signal's path. As the number of wavelengths per fiber increases, and the data rate carried by the wavelengths increases, the amount of electronic processing required in each node grows tremendously. This necessitates the use of very large and expensive electronic switches/routers, as well as a very large amount of electronic terminating equipment, which also presents problems with power consumption and heat dissipation.

In current state-of-the-art core networks, this nodal bottleneck has been greatly reduced by the development of 'optical bypass' technology. This technology allows a signal to remain in the optical domain as it traverses a node, rather than undergoing OEO conversion at every node (this is often referred to as OOO technology). Two technological advances enabled this transformation. First, the transmission 'optical reach' was greatly increased, where the optical reach refers to the distance an optical signal can travel before its quality degrades to a level that necessitates regeneration. Extended optical reach is enabled through the use of, for example, Raman amplification, advanced modulation and error correcting schemes, and dispersion compensation techniques.

The second necessary technological advance was the development of nodal processing equipment that could extract a subset of the wavelengths for electronic processing while allowing the remaining wavelengths to remain in the optical domain. Given that typically 60 to 80% of the wavelengths entering a node do not need to be electronically processed there, such optical bypass equipment eliminates a significant amount of electronic terminating equipment and greatly reduces the size of the required electronic switches. At nodes of degree-two, this equipment is referred to as an *optical add drop multiplexer (OADM)*; at higher degree nodes, it is referred to as an *all-optical switch*.

Given that optical bypass technology has been proven commercially to be more scalable in cost, size and power requirements than traditional OEO-based solutions in large core networks, this technology is assumed as the base architecture for the CORONET target network. The specific assumptions for the optical hardware technology are provided in Table 4, where the first three rows of this table cross-reference the parameters specified for each design scenario of Table 3.

A very important parameter in optical-bypass systems is the optical reach. In general, the optical reach decreases as the bit rate of the wavelength increases. The optical reach that is to be assumed in the CORONET program is specified in Table 4 for the 40 Gb/s line rate and the 100 Gb/s line rate. The optical reach is specified in terms of both a transmission distance and the number of optical network elements traversed. Both limits (i.e., distance and number of elements) need to be simultaneously met in all network designs.

Realistic technical challenges of all-optical solutions, such as transients resulting from switching wavelengths on and off, performance monitoring in the presence of all-optical bypass, etc., should be taken into consideration in proposed solutions.

Even with the presence of ultra-long-reach technology, some amount of regeneration (OEO or OOO) will still be required in the CORONET target network. It can be assumed that a regenerator is capable of full wavelength conversion as well.

Table 4. Optical Hardware Technology Assumptions for the Target Optical Network

Parameter or Item	Assumptions
Capacity per Fiber	Specified in Table 3 for each design scenario.
Number of Fiber Pairs per Network Link	Specified in Table 3 for each design scenario.
Bit Rate per Wavelength	As given in Table 3, the bit rate per wavelength is either 40 Gb/s or 100 Gb/s, depending on the scenario.
Regeneration-Free Optical Reach	The optical reach for the 40 Gb/s wavelengths is assumed to be 2000 km and 10 optical network elements (e.g., all-optical switches or OADMs), and for the 100 Gb/s wavelengths is assumed to be 1500 km and 7 optical network elements.
Node Architecture	Each network node is assumed to include an all-optical switch or an OADM (depending on the degree of the node) to add, drop and/or switch the various wavelengths, as well as to provide optical bypass for traffic that does not need to drop at the node. Electronic equipment located at the nodes (e.g., IP routers and auxiliary cross-connects, if any) is to be connected to the network through the drop-side of the all-optical switch or OADM.
Optical Switching Granularity	Wavelengths or wavebands as dictated by the architecture.
Network Reconfigurability	Completely automated without any manual intervention.
Tunability	All Tx/Rx cards and regenerators (OEO or OOO) are assumed to be tunable across all wavelengths.

The assumptions regarding nodal architecture are very important to the overall network architecture. As specified in Table 4, it is assumed that the core optical network element is either an OADM or an all-optical switch, with any electronic equipment at the node (e.g., IP routers and auxiliary cross-connects, if any) being connected to the network through the drop-side of the all-optical switch or OADM. In such an architecture, IP routers, do not form the basis of the switching platform in the core network. Eliminating the need to electronically process IP traffic at every node along its path has been shown to significantly reduce the router size and the overall network cost.

For purposes of evaluation against the metrics of Section 2.5, assume fully flexible optical network elements; for example, an element can be remotely configured to allow any transmit/receive (Tx/Rx) card plugged into it to add/drop to/from any network port on any wavelength. However, it is also necessary to analyze how the proposed solution would perform using less flexible optical network elements that may have lower cost.

At a minimum, the protocols and algorithms devised in this program should be applicable to the all-optical networking technology described above. In addition, proposers may consider alternative technologies, as long as they are equally (or more) scalable in cost, size and power, for *both* transport and switching. For example, an OEO solution based on highly integrated photonic circuits is a potential alternative technology, but details have to be provided regarding the mechanism of wavelength switching and wavelength bypass at the network nodes. The realism of designs based on such alternative technologies will be crucial in judging the credibility of the proposal. Performers need to provide convincing data to substantiate their assumptions regarding alternative technologies.

2.2.2. Technical Challenges

Two of the important technical challenges of the CORONET program are delivering advancements in both network agility and network resiliency, as described below. A third important technical challenge, which is discussed in Section 2.5.3, involves delivering guaranteed-bandwidth IP services with differentiated QoS. The ability of a proposed solution to meet the CORONET goals in these areas will be measured by the metrics in Section 2.5.

2.2.2.1. <u>Network Agility</u>

One of the fundamental requirements of the CORONET program is the ability to support a highly dynamic traffic environment. In current core networks, most traffic can be considered semi-permanent, with holding times on the order of months or longer. Provisioning new traffic typically takes weeks, with the necessary equipment (e.g., Tx/Rx cards) needed to support a new demand being manually deployed as needed. However, with the emergence of applications such as on-demand computing, where a large data stream is required for a relatively short period of time, the network must be more agile to deliver bandwidth to where it is needed at a given time. The process of connection setup and tear-down needs to be automated so that it can be effected economically and rapidly.

For some applications, such as periodic downloads of large amounts of satellite imagery data, the need for bandwidth is known in advance, and the allocation of resources can be scheduled accordingly. In some situations, such as distributed computation across a set of supercomputers, the need for bandwidth is highly dynamic and unpredictable. Allocating the necessary resources for such services will require a very rapid response time from the network.

The implications and challenges of a highly dynamic network are significant. For example, the control/management plane needs to have fast advanced provisioning capabilities to rapidly establish and tear down connections and flows as required. Up-to-date status of network resources needs to be known by any systems involved with provisioning. Any control protocols used need to converge rapidly and need to be robust with respect to stale state information. Algorithms are needed to rapidly and efficiently route and assign wavelengths to new demands, with the constraint that all required resources for the new demand (e.g., regenerator cards) already be deployed. Automated inventory management is needed to track where equipment has been pre-deployed. Cross-layer capacity management is required to ensure that dynamic IP services are aggregated efficiently in the optical layer. These are just some of the challenges that performers are expected to address as part of successfully supporting the highly dynamic requirements of the CORONET program.

For the purposes of the CORONET program, four classes of traffic have been enumerated with regard to service setup: very fast, fast, scheduled, and semi-permanent. The specific requirements with regard to each class (e.g., setup time, holding time), and the percentage of aggregate network demand in each class, are included in Table 5 in Section 2.5.1.

2.2.2.2. Network Resiliency

Core networks have traditionally provided protection against single link or node failures through the use of dedicated protection resources (throughout this PIP, the terms protection and restoration are used interchangeably). However, although the connections affected by the failure may be rapidly re-established, the failure itself may not be repaired for several hours or even longer. In this repair interval, the network may be subject to additional failures. A global network is especially vulnerable to multiple failures due to its large geographic extent and sometimes-hostile terrain. In the target CORONET network, which potentially will carry mission critical traffic related to the security and defense of the nation, the availability requirements are such that protection against a second, or even a third, failure must be provided.

Not all demands, of course, will require this level of protection. The protection requirements of connections and flows in a core network are typically heterogeneous. For purposes of the CORONET program, three protection classes have been enumerated: protection from a single link or node failure; protection from up to two failures; and protection from up to three failures. (In the classes where protection from multiple failures is required, it is assumed that at most one of the failures is a node failure.) The percentage of aggregate network demand assigned to each protection class is specified, with associated metrics, in Table 6 in Section 2.5.2.

As a general principle, if a failure occurs that does not bring down a particular demand, then that demand should not suffer a hit due to that failure.

Performers are encouraged to develop novel protection schemes that provide rapid restoration, while being efficient with respect to required spare capacity. The specific metrics regarding restoration times and capacity are specified in Section 2.5.2. It is likely that some type of shared mesh protection will be needed in order to meet the capacity metrics. Protection schemes could be, for example, based on shared spare capacity being cross-connected in the electronic layer in response to the failure. As another example, protection could be optical-layer based, relying on all-optical switches to configure the protection paths as needed.

For purposes of evaluation against the restoration metrics of Section 2.5.2, it can be assumed that there is enough time between successive failures such that reallocation of protection resources can occur based on the location of the initial failure(s). However, performers should also analyze how their proposed schemes perform in the event of multiple failures with simultaneous onset times.

Adding to the challenge is the highly dynamic network environment. The protection scheme must remain efficient even as connections and flows enter and leave the network.

Furthermore, the optical bypass environment can be challenging for error recovery because each demand is not electronically processed at every node along its path. The CORONET program encompasses the development of advanced schemes in the areas of performance monitoring, error detection, and fault isolation.

2.3. THRUST AREAS

In order to meet the CORONET objectives, DARPA seeks proposals to design, develop, demonstrate, and evaluate a core network based on novel network architecture and software. Proposers are encouraged to identify, articulate, and exploit a comprehensive approach that will support the advanced network requirements of the target network. Depth of understanding should be evidenced in part through discussion, analysis, and mitigation of proposal-specific risks.

Three major thrust areas encompassed by the CORONET program are discussed below: architecture, protocols/algorithms, and data, control and management planes. This PIP discusses these thrust areas to illuminate the type of thinking that is expected in developing proposals and statements of work, rather than as a prescription. The discussion identifies some, but by no means all, potential avenues that could be pursued in order to achieve the advanced network services required of the target network.

DARPA recognizes that each proposer's understanding of the problem will lead to a unique expression of the solution and a unique statement of work. It is incumbent on each proposer to identify interim milestones and the objective criteria whereby DARPA can determine whether the milestone has been met, recognizing that DARPA desires regular visibility into progress as well as thorough technical documentation. In addition to providing DARPA with interim evaluation criteria, these milestones and objective criteria will be a basis upon which the proposer's depth of understanding of the problem will be evaluated.

2.3.1. Architecture

Networks have evolved rapidly over the past decade, taking advantage of the favorable scaling properties of optical technology (e.g., optical amplifiers and all-optical switches). This trend has not been exhibited to the same degree in core IP routers. The net effect is that as networks continue to grow in capacity, the relative cost of the electronics, compared to the total network cost, continues to increase. The trend of optics to scale better than electronics suggests that optics should play a greater role in architecting future networks as a means of driving cost out of the network.

In the architecture of today's typical core networks, all IP traffic entering the network is processed by an IP router at the edge/core interface, where it is aggregated and packed into wavelengths that are delivered to the core optical layer. As the level of traffic increases, the size of the routers grows accordingly (even when optical-bypass technology is used in the core at intermediate nodes). Maintaining this IP-aggregation architecture in next-generation networks with a ten-fold increase in demand may not be economically feasible.

Optical aggregation is potentially a more scalable approach, where optics rather than electronics is used to map the IP traffic onto the optical layer. Optical aggregation schemes, however, typically rely on scheduling protocols or need to deal with contention in the network. Such schemes may not be feasible in the geographically distributed core network. Rather, they may be more appropriate for the edge network. The relative small size of the edge network (both in geography, number of nodes, and aggregate demand) may allow the use of architectures/protocols that would be impractical in the core. Optical aggregation at the edge is a possible architecture to pursue to reduce the overall network cost, as well as potentially reduce the level of churn in the core network. Other candidate edge-aggregation technologies include Ethernet, Synchronous Optical Network (SONET), G.709 Standard, and Multi-Protocol Label Switching (MPLS). All of these are compatible with IP, but offer simpler data plane forwarding and hence potentially better scalability.

Another possible architectural paradigm is to partition the core network itself into a multi-tier hierarchy; e.g., a two-tier hierarchy consisting of a dense fine-granularity tier and a sparse, express tier.

Another interesting architectural challenge is that of the core switching architecture. As specified in Section 2.2.1.4, all-optical switches are assumed to be the core switching platform for CORONET. However, with an order-of-magnitude increase in traffic, it is not readily obvious whether all-optical switches that operate on individual wavelengths will be economically feasible. The switches may need to operate on wavebands, where groups of wavelengths are switched as a unit. Utilizing a coarser granularity switch will reduce its size and cost, although it has challenges of its own. For example, waveband grooming is likely required, where wavelengths are swapped between wavebands. There is also the architectural issue of whether the waveband size should be fixed throughout the network or be allowed to vary depending on the traffic conditions. Heterogeneous sized wavebands may be more efficient but they are likely more difficult to manage.

The above list of architectural issues and challenges is not comprehensive, and performers are encouraged to meet the objectives of the program by exploiting whatever opportunities they deem most promising. The proposal should include a schedule of milestone reviews during Phase 1 of the program to document and justify the choice of architecture, including a discussion of the long-term limitations and advantages of the chosen architecture. Importantly, the proposal should articulate the objective criteria upon which the chosen architecture would be evaluated by DARPA at that milestone date. The choice of those criteria would also serve as an indicator of problem understanding in the evaluation of the proposal.

2.3.2. Protocols and Algorithms

A host of protocols and algorithms are needed to achieve the advanced network services of the CORONET program. Research on the algorithms and protocols is not restricted to the optical layer. Performers may consider a cross-layer approach where optimization across the IP and optical layers is jointly considered, as long as the results of such research are directly related to the goals of the CORONET program.

One of the biggest challenges in the CORONET program is the requirement to establish some connections in a fraction of a second. Intelligent algorithms will be needed to manage capacity such that bandwidth can be delivered where needed with a very small blocking probability. Capacity management must be efficient; i.e., schemes that rely on setting aside a very large portion of the network capacity in anticipation of future demands are not desirable. Routing demands along very circuitous paths, when better paths exist, is not desirable either. The proposed schemes for automated provisioning must also be cost effective. When a demand is established along a particular path, a small number of regenerator cards may be needed. The demand can only be established along the path if the regenerator cards are pre-deployed at the desired locations. This poses a tradeoff in the cost of pre-deploying equipment versus the ability to, with very high chance of success, provision a demand along a desirable path. Pre-deploying more equipment generally provides more options for routing traffic; however, it clearly comes at a cost. Algorithms for strategically pre-deploying equipment and provisioning demands accordingly will likely play an important role in a highly dynamic network environment. Proposers should provide a suitable metric for assessing the economy of their proposed solution.

This same type of tradeoff may exist with the development of schemes for protection. Developing schemes that are both capacity efficient and economical with respect to the need for pre-deployed equipment is an important goal. Proposers should provide a suitable metric for assessing the economy of their proposed solution.

Another research direction is the optimization of protection across layers. Rather than focusing strictly on the optical layer, performers are encouraged to consider the integration of protection across both the IP and optical layers. Restoration across layers presents many challenges in coordinating the timing of the response of each layer.

Another possible research direction involves scheduled services. Although these services are dynamic, having advance knowledge of when they need to be delivered should allow more efficient use of resources (including protection resources) than service requests that must be fulfilled immediately.

The above protocol/algorithmic issues and challenges are not comprehensive and performers are encouraged to meet the objectives of the program by exploiting whatever opportunities they deem most promising. Any protocols and algorithms that are developed would need to be analyzed, quantified, and demonstrated in order to provide a basis for confidence that they will achieve the program metrics. Proposers should articulate at least one milestone during the first phase in which they document such analyses along with their basis of confidence, whether derived via analysis or simulations. The analysis should consider factors such as implementation complexity, convergence rate, scalability, robustness and security.

The following discussion on the data, control and management planes is closely related to the above discussions of the architecture (Section 2.3.1) and the protocols and algorithms (Section 2.3.2).

2.3.3. Data, Control and Management Planes

A useful way to view a communications network is via the three-plane model, i.e., the data, control and management planes. The data plane is directly involved with the forwarding of packets and bit streams. Most of the above discussion on architecture (Section 2.3.1) is about the data plane. The control and management planes form the mechanisms by which network operations are performed (i.e., fault management, configuration management, accounting, performance monitoring and security). Typically, the control plane takes on a subset of network operations functionality and implements them in a distributed manner across the network, while the management plane is usually oriented along a centralized approach. The protocols and algorithms discussed in Section 2.3.2 above, e.g., for fast service setup and fast restoration, reside mostly in the control plane. An important function of the management plane is to monitor and manage the control plane and to restore it from failures in software or hardware. An important part of this program is the proper division of functions between the control and management planes to enable the reliable and secure delivery of the advanced types of network services discussed in this program. Another important part of the program is the interaction between the control-plane protocols and the data-plane protocols, as well as the horizontal integration of the control planes, e.g., across regional and core networks, as well as vertical integration, e.g., the interplay of the IP and optical control planes.

Security, whether related to the data, control or management planes, is an important thrust of the CORONET program. Security systems include, but are not necessarily limited to, encryption, confidentiality, authentication and authorization techniques, intrusion detection, protection against unauthorized attempts for traffic pattern detection, and protection against denial of service attacks. While the implementation of comprehensive security in the data, control and management planes is not necessary for this program, at least in Phase 1, effects of security systems and measures should be taken into account in architecting the data, control and management planes as well as in assessing the performance of the system via analysis, simulations or emulations. Moreover, software hooks for security systems should be maintained in the code to be developed in Phase 2 of the program. When possible, performers are encouraged to consider emerging standards on security systems and measures rather than invent new proprietary security techniques.

2.4. PROGRAM PHASES

The program has two phases described as follows:

2.4.1. Phase 1

In Phase 1, performers must design the overall network architecture and any necessary protocols and algorithms to meet the goals of the target network as outlined in this PIP. Performers must develop appropriate-level-fidelity models of the overall network system (including the data, control and management planes), and must deliver valid and verifiable simulation results that meet or exceed the threshold program metrics in Section 2.5. (See also Section 2.6.1 for an explanation of the required technical deliverables at the end of Phase 1 of the program.) The period of performance for Phase 1 is 18 months. DARPA expects that Phase 1 of the program is likely to be limited to fundamental research, and thus not subject to U.S. export control restrictions; however, proposers are responsible at all times for complying with all provisions of U.S. law, including those applicable to export control. See Section 3.8 “Export Control.”

2.4.2. Phase 2

In the proposal responding to this BAA, performers are required to submit a high-level proposal for Phase 2 of the program, with a rough-order-of-magnitude (ROM) estimate of the cost, which will be used by the Government for budgeting purposes only. The goals of Phase 2 of the program are to develop, integrate, demonstrate and test network control and management software implementing the architecture, protocols and algorithms devised in Phase 1. An important objective of Phase 2 is to make the results of the CORONET program attractive for transition to commercial carriers. Thus, a credible commercial transition plan needs to be included in the high-level proposal for Phase 2. An important consideration in Phase 2 is to ensure the proposed solution addresses realistic challenges of network deployment. For example, the proposed solution should consider: the evolution of optical network elements over the next several years; hybrid line-rate systems; horizontal interactions of the control plane across heterogeneous core networks; etc.

The period of performance of Phase 2 is to be specified by each proposer. Phase 2 of the program may be subject to U.S. export-control control laws and regulations.

Three months prior to the end of Phase 1, performers will be required to submit a full (technical and cost) proposal for Phase 2 of the program, whose high-level goals are defined above. DARPA will select the performers for Phase 2 based on the successful demonstration of the Phase 1 goals and metrics as well as the comprehensiveness and realism of the Phase 2 proposal, its overall adherence to the CORONET program goals, and the credibility of its commercial transition plan.

New teaming arrangements can be proposed for Phase 2 of the program in order to enhance its successful completion and commercial transition. Moreover, to enhance the cost effectiveness of Phase 2 of the program, some parts of the proposed development and test plans should be included in the proposal as fully costed *Options*. This is particularly important if some development is proposed to implement new requirements that go beyond those presented in this PIP. These program *Options* may or may not be exercised by the Government at the time of selection of the performers for Phase 2.

DARPA anticipates updating the requirements for Phase 2 of this BAA in the latter half of Phase 1 to give the performers sufficient time to write their full Phase 2 proposals.

2.5. PROGRAM METRICS

In order for the Government to evaluate the effectiveness of a proposed solution in achieving the stated program objectives, proposers should note that the Government hereby promulgates the following program metrics that may serve as the basis for determining whether satisfactory progress is being made to warrant continued funding of the program. Although the following program metrics are specified, proposers should note that the government has identified these goals with the intention of bounding the scope of effort, while affording the maximum flexibility, creativity, and innovation in proposing solutions to the stated problem.

The metrics presented below are merely threshold metrics, and results that exceed the specified program metrics are highly desirable. Additional metrics of performance exist, and performers are encouraged to identify other characteristics that they propose to quantitatively evaluate over and above the Government's independent test and evaluation.

All metrics in this section apply to each of the four design scenarios that were given in Table 3. The metrics must be met in the simulation and modeling of Phase 1 as well as in the more comprehensive testing (e.g., network emulation) that will potentially be made in Phase 2.

Any results obtained via simulation should be based on a sufficient number of runs to produce reasonable confidence intervals. In the simulations, state-of-the-art CPUs should be assumed to be deployed at the network nodes to perform network control and management duties. Moreover, as mentioned in Section 2.3.3, overhead due to security measures should be taken into account in the simulations.

2.5.1. Service Setup Metrics

The CORONET target network is highly dynamic and will support four classes of service setup, ranging from Very Fast to Semi-Permanent. The characteristics of each class (e.g., setup time, holding time), as well as the percentage of aggregate network demand in each class, are specified in Table 5. The traffic breakdown by setup class is specified as a function of the parameter *S*. As defined in Table 2, *S* represents the fraction of the aggregate network demand that is IP services; 1- *S* is the fraction that is wavelength services. Demands should be arbitrarily assigned to one of the four setup classes in accordance with the specifications of Table 5.

Note that best-effort IP services are not assigned to any of the service setup classes. These services can be considered as always-present, albeit highly variable traffic.

For the Very Fast, Fast, and Scheduled Service classes, the demands must be provisioned using equipment that has already been deployed in the network. For the semi-permanent traffic, it is assumed that the provisioning times are long enough such that equipment can be installed as needed.

Note that traffic in the Very Fast Setup class, which consists solely of wavelength services, is protected against only single failures. Traffic assigned to the other three setup classes is protected as specified in Table 6.

The blocking probability thresholds that must be met for each class are specified in the table. These blocking thresholds apply to each service setup attempt. If a service setup attempt fails, additional attempts can continue to be made until it finally succeeds.

Clearly there is a tradeoff between available network capacity and the blocking probability for dynamic services. The more capacity that is deployed in the network, the lower the expected blocking probability. To ensure solutions are capacity efficient, a limit on the available network capacity is imposed, as will be described in the next section. Thus, schemes that rely on low link load in order to operate effectively will likely not meet this bound.

Table 5. Performance Metrics for Service Setup (applies to each design scenario given in Table 3)

		Service Setup Classes			
		Very Fast Service Setup	**Fast Service Setup**	**Scheduled Service Setup**	**Semi-Permanent**
Traffic Breakdown [1]	**Guaranteed-Bandwidth IP Services**	0%	20% of S	20% of S	20% of S
	Wavelength Services	40% of (1-S)	20% of (1-S)	20% of (1-S)	20% of (1-S)
Assumptions		These three cases represent traffic churn that must be satisfied using pre-deployed equipment			Equipment can be installed as needed
Setup Time		CONUS: ≤ **100 msec** [2] Global: ≤ **250 msec** [2]	≤ **2 sec**	≤ **10 sec**	N/A
Holding Time		1 sec to 1 min	10 sec to 10 hr	1 min to 1 Month	Months
Protection		Single Failures Only	See Table 6		
Blocking Probability		$\mathbf{10^{-3}}$	$\mathbf{10^{-3}}$	$\mathbf{10^{-4}}$	N/A
Network-Capacity Limit		See Table 6			

(1) This represents the fraction of the aggregate network demand that is in the various setup-time classes. The aggregate network demand is provided in Table 3 for four design scenarios. *S* is defined in Table 2. Note that best-effort IP services (which comprise 40% of *S*, as specified in Table 2) are not included here.

(2) The required very fast service-setup times can be defined more generally as: 50 msec + The round-trip fiber transmission delay between the source and destination nodes.

2.5.2. Restoration Metrics

The restoration classes that must be supported in the CORONET target network are specified in Table 6. These classes represent different quality of service levels with respect to protection. For each class, the number and type of network failures (i.e., node or link failure) for which traffic restoration must be provided is specified.

The percentage of aggregate network demand in each restoration class is specified in Table 6, as a function of the protection parameter *P*. In the network studies that are to be performed as part of the CORONET program, *P* ranges from 0% to 5%. The nominal value of *P*, denoted by P_o, is 2.5% (this is the value to be assumed whenever the PIP refers to a nominal network design). Note that best-effort IP services, which represent 40% of the IP services, are not assigned to any of the explicit restoration classes specified in Table 6.

As stated in the previous section, traffic in the Very Fast Setup class is assigned solely to the Single-Failure Restoration class. This traffic may be restored, for example, by issuing another service setup request.

All other demands requiring protection should be arbitrarily assigned to the various restoration classes. It is not acceptable, for example, to assign triple failure restoration to only those demands that are the easiest to protect. Also, demands can only be assigned to a restoration class if it is topologically possible that the requirements can be met. For example, a demand cannot be protected against all triple failures if there are only two diverse paths between the demand endpoints.

In the restoration classes where protection from multiple failures is required, it is assumed that at most one of the failures is a node failure. For purposes of meeting the metrics in this table, it can be assumed that there is enough time between successive failures such that reallocation of protection resources can occur based on the location of the initial failure(s).

Note that when a network node fails, all equipment at that node (optical network elements, IP routers, and other equipment) is assumed to have failed. Moreover, traffic originating at the failed node does not need to be restored. Note also that when a node fails, the control and management functions residing at that node will no longer be available. Also, when a network link between two nodes fails, the control and management traffic between these nodes will be interrupted.

Table 6 addresses only a subset of the failures in a network. Any proposed solution must be robust to a range of failures, e.g., failures of Tx/Rx cards, IP router ports, control plane, etc.

It is required that performers develop restoration schemes that are efficient in their use of protection capacity. Table 6 includes a metric regarding the ratio of the protection capacity to the working capacity. This *spare capacity ratio* is defined further here. Assume *D* represents the set of all active demands in the network. Let *W* be the total capacity allocated to the working paths of the demands in *D*, as measured in wavelength-km. (The resources allocated to best-effort IP services should be included in *W*.) Let *B* be the total allocated backup capacity for the demands assigned to the restoration classes of Table 6, again measured in wavelength-km. (This backup capacity can be dedicated or shared, or a combination of both.) The spare capacity ratio is then defined as: B/W.

Table 6. Performance Metrics for Fast Restoration from Network Failures (applies to each design scenario given in Table 3)

	Restoration Classes		
	Single-Failure Restoration	**Double-Failure Restoration**	**Triple-Failure Restoration**
Description	Capable of restoration from any single network failure (i.e., a link or node failure)	Capable of restoration from up to two network failures, with no more than one node failure (3), (5)	Capable of restoration from up to three network failures, with no more than one node failure (4), (5)
Traffic Breakdown (1)	(1 - 0.4 S) - 5 P	4 P	P
	Where the *Protection Parameter* $P \in [0\%, 5\%]$. The nominal value is $P_o = 2.5\%$		
Restoration Time (2)	*Single-Failure Restoration Time (SFRT):* (2) CONUS: ≤ **100 msec** Global: ≤ **250 msec**	First failure: ***SFRT*** Second failure: ***SFRT* + 50 msec**	First failure: ***SFRT*** Second failure: ***SFRT* + 50 msec** Third failure: ***SFRT* + 100 msec**
Total Spare Protection Capacity Relative to Total Working Capacity	This ratio should be minimized as much as possible for each network design. For the nominal network designs (6), this ratio should **not exceed 75%** for the CONUS-only capacities. (See Section 2.5.2 for details on how this ratio is to be measured)		
Network-Capacity Limit	Ideally, the fast-restoration metrics given above, as well as the service-setup metrics given in Table 5 should be simultaneously met with the link capacities specified in Table 3 for each design scenario. Extra capacity may be added if necessary, but the number of network links requiring one extra fiber pair should be minimized as much as possible. For the nominal network designs (6), **no more than 10%** of the network links should have one extra fiber pair.		

(1) This row gives the fraction of the aggregate network demand requiring the corresponding protection level. This row is specified in terms of *P*, where *P* can take on a range of values, as specified. All traffic in the Very Fast Setup class is assigned to single-failure restoration (see Table 5). All other protected demands should be arbitrarily assignable across the three protection classes, as long as the distribution of traffic falls within the indicated percentages. Best-Effort IP Services, which represent 40% of the IP traffic, are not included in any of the three restoration classes.

(2) *Single-Failure Restoration Time (SFRT)* can be defined more generally as: 50 msec + The round-trip fiber transmission delay between the source and destination nodes.

(3) Only applies to traffic between node pairs with network connectivity ≥ 3.

(4) Only applies to traffic between node pairs with network connectivity ≥ 4.

(5) In the case of multiple failures, assume there is enough time between failures to recompute backup paths.

(6) For each of the design scenarios of Table 3, the nominal network design parameters are S_o and P_o (S_o is defined in Table 2).

(Note that this metric implicitly assumes that working and backup traffic are carried in separate wavelengths. If performers propose a protection scheme that combines working and backup traffic in the same wavelength, they must demonstrate that it provides better performance, and they must define an appropriate alternative metric that captures the ratio of backup spare capacity to working capacity.)

The target CORONET network is highly dynamic, with D continually changing; thus the spare capacity ratio B/W will change as the network evolves. This ratio should be calculated at frequent regular intervals, with a long-term average of this ratio ultimately calculated for each scenario. Any reference below to the spare capacity ratio refers to this long-term average.

For each of the design scenarios of Table 3, performers must provide the following plots as part of Phase 1:

- The spare capacity ratio vs. the parameter P, with the parameter S fixed at S_o
- The spare capacity ratio vs. the parameter S, with the parameter P fixed at P_o
- W vs. the parameter P, with the parameter S fixed at S_o
- W vs. the parameter S, with the parameter P fixed at P_o

It is desirable to further evaluate the results of these designs with regard to the CONUS-only portion of the network. Let W_c be the total working capacity allocated in CONUS, measured in wavelength-km. Similarly, let B_c be the total backup capacity allocated in CONUS. (If a demand is partially routed in CONUS, then only the CONUS portion of the working and backup paths should be included in W_c and B_c, respectively.) For each of the design scenarios, the same four plots listed above must also be provided in Phase 1 using W_c and B_c.

It is desirable that the spare capacity ratio be as low as possible, while meeting the requirements of the various protection classes. To provide a firm metric, Table 6 specifies that for S_o and P_o, the long-term average ratio B_c/W_c should be no greater than 75%. This metric must be met in Phase 1 for all four design scenarios of Table 3.

In addition to capping the percentage of spare capacity, Table 6 also includes a metric regarding the total capacity in the network. As indicated in Table 6, the fast restoration requirements of this section, as well as the service setup requirements of the previous section, should be simultaneously met with the link capacities that are specified in Table 3 for each of the four design scenarios. It is recognized that this may not be possible depending on the demand pattern and the network topology chosen. Thus, *one* extra fiber pair may be added to a subset of the network links, where the number of such links should be minimized as much as possible. Performers must provide the following two plots as part of Phase 1, for each of the four design scenarios of Table 3:

- The number of links with extra capacity vs. the parameter P, with the parameter S fixed at S_o
- The number of links with extra capacity vs. the parameter S, with the parameter P fixed at P_o

To provide a firm metric regarding total network capacity, Table 6 specifies that for the nominal network design (i.e., S_o and P_o), no more than 10% of the network links should have an extra fiber pair. This metric must be met for all four design scenarios of Table 3.

Also, in Phase 1, the following plots must be provided for each design scenario:

- The maximum number of wavelengths ever required on any link vs. the parameter *P*, with the parameter *S* fixed at S_o
- The maximum number of wavelengths ever required on any link vs. the parameter *S*, with the parameter *P* fixed at P_o

2.5.3. IP Service Metrics

The quality of service (QoS) metrics for each of the IP service types is specified in Table 7. Latency is the end-to-end delay across the core network; jitter is the variation in this delay; and packet loss is the fraction of packets that are dropped. The specifications apply to each of the design scenarios that are given in Table 3. The metrics in Table 7 must be met in the simulation and modeling of Phase 1 as well as in the comprehensive testing in Phase 2.

The percentage of aggregate network demand in each QoS class is specified in Table 7 as a function of *S*, where *S* is defined in Table 2. The guaranteed-bandwidth services should be arbitrarily partitioned (independently from data rate) into services that are sensitive to latency and jitter (e.g., interactive video) and services that are sensitive to packet loss (e.g., mission-critical data transfer).

For best-effort IP services, the QoS metrics are not guaranteed for any individual flow; the metrics must be met with regard to long-term averages. For guaranteed-bandwidth services, the QoS metrics must be met for every individual flow.

Note that for purposes of meeting the IP QoS metrics of Table 7, performers can assume that the network is a single administrative domain.

Table 7. QoS Metrics for IP Services (applies to each design scenario given in Table 3)

	Best-Effort IP Services [2]	**Guaranteed-Bandwidth IP Services** [3]	
		Latency- and Jitter-Sensitive Services	**Packet-Loss-Sensitive Services**
Traffic Breakdown [1]	40% of S	30% of S	30% of S
One-Way Latency	CONUS: **100 msec** [4] Global: **250 msec** [4]	CONUS: **50 msec** [5] Global: **125 msec** [5]	**500 msec**
One-Way Jitter [6]	CONUS: **20 msec** Global: **50 msec**	CONUS: **10 msec** Global: **25 msec**	**100 msec**
Packet Loss	**10^{-3}**	**10^{-3}**	**10^{-6}**

(1) This row represents the fraction of the aggregate network demand that is in each QoS category. The aggregate network demand is provided in Table 3 for four design scenarios. *S* is defined in Table 2.

(2) For Best-Effort IP Services, the QoS metrics are statistical averages that are not guaranteed for any individual flow.

(3) For Guaranteed-Bandwidth IP Services, the QoS metrics are to be guaranteed for every individual flow.

(4) In general, this is equal to: 2×(25 msec + the one-way, source-to-destination fiber transmission delay).

(5) In general, this is equal to: 25 msec **+** the one-way, source-to-destination fiber transmission delay.

(6) In general, this is 20% of the one-way latency.

2.6. TECHNICAL DELIVERABLES

Some of the important technical deliverables that are expected as part of the two program phases are highlighted here. Note that this is not meant to be an exhaustive list of all deliverables; it is expected that performers will deliver the necessary technical output to demonstrate a comprehensive solution that meets the goals of the CORONET program.

2.6.1. Phase 1

At a minimum, the following technical items must be delivered at the end of Phase 1 of the program:

- Detailed description of the node architecture, including the hardware architecture of the optical network elements, giving sufficient details to enable assessing the technology being proposed.
- Detailed description of the architecture of the control and management planes, including the fault management, configuration management and performance monitoring functions as well as security concerns.
- A comprehensive description of all protocols and algorithms that have been developed. The analysis should include factors such as implementation complexity, convergence rate, scalability, robustness and security concerns.
- Appropriate-level-fidelity models of any protocols and algorithms. Note that protocol correctness simulations could use a very detailed or high fidelity model, while comprehensive network simulations may use more approximate models to increase simulation speed. Models used for simulation must clearly show how such ‘environmental’ factors as processor speed, processor loading, limitations on control plane bandwidth, etc. are taken into account.
- Detailed description of the mechanism(s) used to model a cross-layer IP over WDM environment, and the methodology used to ensure and measure the IP QoS metrics of Table 7.
- Illustration of the global network topology used in the Phase 1 simulations, as well as a representation of the topology (i.e., cities, links, link distances) in MS Excel-readable format.
- Snapshot of the traffic distribution for a nominal network design for each of the four scenarios in Table 3. The demand endpoints, service type, data rate, and protection level should be provided in MS Excel-readable format. The program encompasses traffic churn, so clearly the distribution changes over time. It is necessary to submit only a snapshot of the traffic distribution at some instant in time, once the network has reached a steady state. Also, provide a description of the mechanism used to model the traffic churn in the simulations.
- Verifiable simulation results that demonstrate that the proposed solution meets all of the program metrics specified in Section 2.5 for each design scenario. By verifiable, it is meant that example network and models could be run, with assistance from proposers, on an OPNET® core simulation platform at a facility specified by the program office. In the simulations, state-of-the-art CPUs should be assumed to be deployed at the network

nodes to perform network control and management duties. Moreover, as mentioned in Section 2.3.3, overhead due to security measures should be taken into account in the simulations. (See Section 2.7 below for a high-level discussion of independent test and evaluation.)

- Plots of spare capacity ratio and working capacity as a function of *S* and *P* for each design scenario (see Section 2.5.2 for more details).
- Plots of overall network capacity measures as a function of *S* and *P* for each design scenario (see Section 2.5.2 for more details on these measures).
- For the nominal network design (S_o and P_o), statistics on:
  - average (weighted by bit rate) and maximum end-to-end working path length for all demands
  - average and maximum link utilization in terms of wavelengths
  - average wavelength fill rate (for the wavelengths carrying IP services)
  - average wavelength bypass of the IP routers across the network (i.e., the percentage of wavelengths that enter nodes from the network side that do not enter the IP routers)
  - average wavelength bypass of nodes across the network (i.e., the percentage of wavelengths that enter nodes from the network side that directly pass through without being dropped or regenerated (either OEO or OOO))
  - average size of an IP router across the core network, and maximum size of an IP router
  - any other statistics relevant to the particular proposed solution that can shed light on the economics and the technological feasibility of the solution
- Quantitative analysis that demonstrates the proposed solution exhibits scalability in cost, size, and power requirements.
- Discussion regarding support of:
  - Multicast for both IP and wavelength services (Section 2.2.1.2)
  - Hybrid line-rate systems (Section 2.2.1.3)
  - Evolution of optical network elements (Section 2.2.1.4)
  - Multiple failures with simultaneous onset (Section 2.2.2.2)
- Full (technical and cost) proposal for Phase 2 (submitted three months before the end of Phase 1) detailing proposed development program, milestones, metrics and test plan that would make the results of the CORONET program attractive for commercial transition (see Section 2.4.2 for more details).

2.6.2. Phase 2

The scope and technical deliverables for Phase 2 of the program will be defined in the full proposal for Phase 2 to be submitted by the performers three month prior to the end of Phase 1. DARPA also anticipates issuing some of the Phase 2 requirements and technical deliverables when this BAA is updated during the latter half of Phase 1.

2.7. INDEPENDENT TEST AND EVALUATION DURING PHASE 1 OF THE PROGRAM

This program will utilize a contractor team at an independent test lab for verification and evaluation of simulations. The test lab will be involved with the following items:

- Working with performers to ensure the realism of simulation models. For example, this would include making sure that all performers take into account various limitations present in real networks, such as control/management plane processor speed, control/management plane bandwidth limitations, delays in real networks, etc.
- Working with performers to establish a common discrete event simulation system and level of fidelity appropriate for the system or subsystem being simulated. For example, protocol verification may use a much higher level of fidelity than statistical simulations aimed at measuring blocking probabilities.
- Work with performers to select a representative set of systems or subsystems and scenarios to be subjected to independent test and verification.
- Independently test and verify performers' models in designated scenarios, using performers' code and simulation configurations running within the simulation system at the independent test lab. In order to facilitate analysis, assessment, and follow-on transition of software, simulations must be constructed around an OPNET® core simulation and utilize OPNET® protocol and technology modules to the extent possible.

3. GENERAL INFORMATION

3.1. ELIGIBILITY

This BAA solicits proposals from all interested and qualified sources. Foreign participants and/or individuals may participate to the extent that such participants comply with any necessary Non-Disclosure Agreements, Security Regulations, Export Laws, and other governing statutes applicable under the circumstances.

3.2. LIMITATIONS ON OTHER TRANSACTION FOR PROTOTYPE PROJECTS

Proposers are advised that an Other Transaction for Prototype Agreement will only be awarded if there is:

1. At least one nontraditional defense contractor participating to a significant extent in the prototype project, or
2. No nontraditional defense contractor is participating to a significant extent in the prototype project, but at least one of the following circumstances exists:
    a. At least one third of the total cost of the prototype project is to be paid out of funds provided by the parties to the transaction other than the federal Government. The cost share should generally consist of labor, materials, equipment, and facilities costs (including allocable indirect costs).
    b. Exceptional circumstances justify the use of a transaction that provides for innovative business arrangements or structures that would not be feasible or appropriate under a procurement contract.

Although use of one of these options is required to use an Other Transaction for Prototype agreement as the procurement vehicle, no single option is encouraged or desired over the others. However, DARPA has not used the exceptional circumstances justification for the BAA process.

NOTE: For purposes of determining whether or not a participant may be classified as a nontraditional defense contractor and whether or not such participation is determined to be participating to a significant extent in the prototype project, the following definitions are applicable:

"Nontraditional defense contractor" means a business unit that has not, for a period of at least one year prior to the date of the OT agreement, entered into or performed on:

1. any contract that is subject to full coverage under the cost accounting standards prescribed pursuant to section 26 of the Office of Federal Procurement Policy Act (41 U.S.C. 422) and the regulations implementing such section; or

2. any other contract in excess of $500,000 to carry out prototype projects or to perform basic, applied, or advanced research projects for a Federal agency that is subject to the Federal Acquisition Regulation.

"Participating to a significant extent in the prototype project" means that the nontraditional defense contractor is supplying a new key technology or product, is accomplishing a significant amount of the effort wherein the role played is more than a nominal or token role in the research effort, or in some other way plays a significant part in causing a material reduction in the cost or schedule of the effort or an increase in performance of the prototype in question.

NOTE: Proposers are cautioned that if they are classified as a traditional defense contractor, and propose the use of an OT for Prototype Agreement, the government will require submittal of both a cost proposal under the guidelines of the FAR/DFARS, and a cost proposal under the proposed OT for Prototype Agreement, so that an evaluation may be made with respect to the cost tradeoffs applicable under both situations. The government reserves the right to negotiate either a FAR based procurement contract, or Other Transaction for Prototype Agreement as it deems is warranted under the circumstances.

3.3. PROCUREMENT INTEGRITY, STANDARDS OF CONDUCT, ETHICAL CONSIDERATIONS

Certain post-employment restrictions on former federal officers and employees may exist, including special Government employees (including but not limited to Section 207 of Title 18, United States Code, the Procurement Integrity Act, 41 U.S.C. 423, and FAR 3.104). If a prospective proposer believes that a conflict of interest exists, the situation should be raised to the DARPA Contracting Officer before time and effort are expended in preparing a proposal. All proposers and proposed sub-contractors must therefore affirm whether they are providing scientific, engineering, and technical assistance (SETA) or similar support to any DARPA technical office(s) through an active contract or subcontract. All affirmations must state which office(s) the proposer supports and identify the prime contract numbers. Affirmations shall be furnished at the time of proposal submission. All facts relevant to the existence or potential existence of organizational conflicts of interest (FAR 9.5.) must be disclosed. The disclosure shall include a description of the action the proposer has taken or proposes to take to avoid, neutralize, or mitigate such conflict.

3.3.1. Human Use In Research

Proposals selected for funding are required to comply with provisions of the Common Rule (32 CFR 219) on the protection of human subjects in research (http://www.dtic.mil/biosys/downloads/32cfr219.pdf) and the Department of Defense Directive 3216.2 (http://www.dtic.mil/whs/directives/corres/html2/d32162x.htm). All proposals that involve the use of human subjects are required to include documentation of their ability to follow Federal guidelines for the protection of human subjects. This includes, but is not limited to, protocol approval mechanisms, approved Institutional Review Boards, and Federal Wide Assurances. These requirements are based on expected human use issues some time during the entire length of the proposed effort.

For proposals involving "greater than minimal risk" to human subjects within the first year of the project, performers must provide evidence of protocol submission to a federally approved IRB at the time of final proposal submission to DARPA. For proposals that are forecasted to involve "greater than minimal risk" after the first year, a discussion on how and when the proposer will comply with submission to a federally approved IRB needs to be provided in the submission. More information on applicable federal regulations can be found at the Department of Health and Human Services - Office of Human Research Protections website (http://www.dhhs.gov/ohrp/).

3.4. INTELLECTUAL PROPERTY

3.4.1. Procurement Contract Proposers

3.4.1.1. <u>Noncommercial Items: (Technical Data and Computer Software)</u>
Proposers responding to this BAA requesting a procurement contract to be issued under the FAR/DFARS, shall identify all noncommercial technical data, and noncommercial computer software that it plans to generate, develop, and/or deliver under any proposed award instrument in which the Government will acquire less than unlimited rights, and to assert specific restrictions on those deliverables. Proposers shall follow the format under DFARS 252.227-7017 for this stated purpose. In the event that proposers do not submit the list, the Government will assume that it automatically has “unlimited rights” to all noncommercial technical data and noncommercial computer software generated, developed, and/or delivered under any award instrument, unless it is substantiated that development of the noncommercial technical data and noncommercial computer software occurred with mixed funding. If mixed funding is anticipated in the development of noncommercial technical data, and noncommercial computer software generated, developed, and/or delivered under any award instrument, then proposers should identify the data and software in question, as subject to Government Purpose Rights (GPR). In accordance with DFARS 252.227-7013 Rights in Technical Data - Noncommercial Items, and DFARS 252.227-7014 Rights in Noncommercial Computer Software and Noncommercial Computer Software Documentation, the Government will automatically assume that any such GPR restriction is limited to a period of five (5) years in accordance with the applicable DFARS clauses, at which time the Government will acquire “unlimited rights” unless the parties agree otherwise. Proposers are admonished that the Government may use the list during the source selection evaluation process to evaluate the impact of any identified restrictions, and may request additional information from the proposer, as may be necessary, to evaluate the proposer’s assertions. If no restrictions are intended, then the proposer should state “NONE.”

A sample list for complying with this request is as follows:

NONCOMMERCIAL			
Technical Data Computer Software To be Furnished With Restrictions	Basis for Assertion	Asserted Rights Category	Name of Person Asserting Restrictions
(LIST)	(LIST)	(LIST)	(LIST)

3.4.1.2. Commercial Items: (Technical Data and Computer Software)
Proposers responding to this BAA requesting a procurement contract to be issued under the FAR/DFARS, shall identify all commercial technical data, and commercial computer software that may be embedded in any noncommercial deliverables contemplated under the research effort, along with any applicable restrictions on the Government's use of such commercial technical data and/or commercial computer software. In the event that proposers do not submit the list, the Government will assume that there are no restrictions on the Government's use of such commercial items. The Government may use the list during the source selection evaluation process to evaluate the impact of any identified restrictions, and may request additional information from the proposer, as may be necessary, to evaluate the proposer's assertions. If no restrictions are intended, then the proposer should state "NONE."

A sample list for complying with this request is as follows:

COMMERCIAL			
Technical Data Computer Software To be Furnished With Restrictions	Basis for Assertion	Asserted Rights Category	Name of Person Asserting Restrictions
(LIST)	(LIST)	(LIST)	(LIST)

3.4.2. NonProcurement Contract Proposers - Noncommercial and Commercial Items: (Technical Data and Computer Software)

Proposers responding to this BAA requesting a Technology Investment Agreement or Other Transaction for Prototype shall follow the applicable rules and regulations governing these various award instruments, but in all cases should appropriately identify any potential restrictions on the Governments use of any Intellectual Property contemplated under those award instruments in question. This includes both Noncommercial Items and Commercial Items. Although not required, proposers may use a format similar to that described in Paragraphs 3.4.1 and 3.4.2 herein. The Government may use the list during the source selection evaluation process to evaluate the impact of any identified restrictions, and may request additional information from the proposer, as may be necessary, to evaluate the proposer's assertions. If no restrictions are intended, then the proposer should state "NONE."

3.4.3. All Proposers – Patents

Please include documentation proving your ownership of or possession of appropriate licensing rights to all patented inventions (or inventions for which a patent application has been filed) that will be utilized under your proposal for the DARPA program. If a patent application has been filed for an invention that your proposal utilizes, but the application has not yet been made publicly available and contains proprietary information, you may provide only the patent number, inventor name(s), assignee names (if any), filing date, filing date of any related provisional application, and a summary of the patent title, together with either: 1) a representation that you own the invention, or 2) proof of possession of appropriate licensing rights in the invention.

3.4.4. All Proposals - Intellectual Property Representations

Please provide a good faith representation that you either own or possess appropriate licensing rights to all intellectual property that will be utilized under your proposal for the DARPA program**.** If you are unable to make such a representation concerning non-patent related intellectual property, please provide a listing of the intellectual property to which you do not have needed rights, and provide a detailed explanation concerning how and when you plan to obtain these rights.

3.5. REPORTING REQUIREMENTS

The number and types of reports will be specified in the award document, but will include as a minimum quarterly R&D and financial status reports (see sample at attachment A). The reports shall be prepared and submitted in accordance with the procedures contained in the award document and mutually agreed on before award. A Final Report that summarizes the project and tasks will be required at the conclusion of the performance period for the award, notwithstanding the fact that the research may be continued under a follow-on vehicle. Each performer will also be required to submit periodic reports on invention disclosure, election of title, and filing of patent applications.

3.6. REQUIRED REVIEW AND INTERCHANGE MEETINGS

The following schedule of technical and programmatic meetings will be included in the program schedule:

- Program Kick-off meeting at the performer’s facilities within one month after program award.
- Periodic program (milestones, schedule, budget) and technical performance reviews at performer’s facilities every four months, beginning in month 5 of the program.
- Critical Design review at the performer’s facilities at the month 9 meeting.
- Discrete Event Simulation Test Plan Review at the performer’s facilities at the month 13 meeting.
- Final demonstration of simulation results at performer’s facilities in month 17.
- Final Phase 1 Program Review and Phase 2 briefing at performer’s facilities in month 18.

3.7. SUBCONTRACTING

Pursuant to Section 8(d) of the Small Business Act (15 U.S.C. 637(d)), it is the policy of the Government to enable small business and small disadvantaged business concerns to be considered fairly as subcontractors to contractors performing work or rendering services as prime contractors or subcontractors under Government contracts, and to assure that prime contractors and subcontractors carry out this policy. Each proposer who submits a contract proposal and includes subcontractors is required to submit a subcontracting plan IAW FAR 19.702(a) (1) and (2) should do so with their proposal. The plan format is outlined in FAR 19.704.

3.8. EXPORT LICENSES

The following, or a similar, provision will be incorporated into any resultant contract:

(1) The Contractor agrees to comply with the requirements of all applicable U.S. export control laws and regulations, including, but not limited to, the Export Administration Regulations (EAR) 15 CFR 730 et seq. and the International Traffic in Arms Regulations (ITAR) 22 CFR 120 et seq.. The Contractor agrees it is solely responsible for determining the applicability of such laws and regulations to its performance of this Contract and for ensuring it remains in compliance with applicable laws and regulations at all times.

(2) The Contractor shall be solely responsible for obtaining all required export licenses before performing under this Contract. The Contractor acknowledges its responsibilities may include obtaining appropriate licenses for deemed exports if the Contractor plans to use foreign nationals in the performance of this Contract, including for work that may be performed on-site at government installations located in the United States.

(3) The Contractor shall be responsible for all regulatory record keeping requirements associated with the use of export licenses and license exemptions and exceptions.

(4) The Contractor agrees to include the provisions of this clause, appropriately modified to reflect the parties, in all of its subcontracts pertaining to work performed under this Contract.

3.9. PUBLIC RELEASE OR DISSEMINATION OF INFORMATION

The following provision will be incorporated into any resultant contract:

(a) There shall be no dissemination or publication, except within and between the Contractor and any subcontractors, of information developed under this contract or contained in the reports to be furnished pursuant to this contract without prior written approval of the Contracting Officer Representative (COR). All technical reports will be given proper review by appropriate authority to determine which Distribution Statement is to be applied prior to the initial distribution of these reports by the Contractor. Papers resulting from unclassified contracted fundamental research are exempt from prepublication controls and this review requirement, pursuant to DoD Instruction 5230.27 dated October 6, 1987.

(b) When submitting material for clearance for open publication, the Contractor must furnish DARPA Technical Information Officer, 3701 North Fairfax Drive, Arlington VA 22203-1714, telephone (703) 526-4163 with five copies and allow four weeks for processing. Viewgraph presentations must be accompanied with a written text. Whenever a paper is to be presented at a meeting, the Contractor must indicate the exact dates of the meeting or the Contractor’s date deadline for submitting the material.

3.10. AWARD ADMINISTRATION INFORMATION

(1) Central Contractor Registration. Selected offerors not already registered in the Central Contractor Registry (CCR) will be required to register in CCR prior to any award under this BAA. Information on CCR registration is available at http://www.ccr.gov.

(2) Representations and Certifications. In accordance with Federal Acquisition Regulation 4.1201, prospective proposers shall complete electronic annual representations and certifications at http://orca.bpn.gov.

4. PROPOSAL PREPARATION

4.1. GENERAL GUIDANCE

All proposals submitted must follow the instructions in this Proposer Information Pamphlet (PIP) and include only the information requested to avoid delays in evaluation or disqualification. It is anticipated that within 30 days of completing the evaluation, proposers will be notified that: 1) its proposal has been selected for negotiation, or 2) its proposal has not been selected. Proposals not selected will be destroyed; however, a copy of non-selected proposals may be retained and filed.

4.1.1. Restrictive Markings on Proposals

All proposals should clearly indicate limitations on the disclosure of their contents. Proposers who include in their proposals data that they do not want disclosed to the public for any purpose, or used by the Government except for evaluation purposes, shall-

(1) Mark the title page with the following legend:

> This proposal includes data that shall not be disclosed outside the Government and shall not be duplicated, used, or disclosed-in whole or in part-for any purpose other than to evaluate this proposal. If, however, a contract is awarded to this proposer as a result of, or in connection with, the submission of this data, the Government shall have the right to duplicate, use, or disclose the data to the extent provided in the resulting contract. This restriction does not limit the Government's right to use information contained in this data if it is obtained from another source without restriction. The data subject to this restriction are contained in sheets [*insert numbers or other identification of sheets*]; and

(2) Mark each sheet of data it wishes to restrict with the following legend:

> Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this proposal.

Markings like "Company Confidential" or other phrases that may be confused with national security classifications shall be avoided.

4.1.2. Confidentiality

It is the policy of DARPA to treat all proposals as competitive information and to disclose their contents only for the purpose of evaluation. No proposals will be returned. The original of each proposal received will be retained at DARPA and all other copies of non-selected proposals destroyed. Documentation related to the source selection process will be marked SOURCE SELECTION INFORMATION – SEE FAR 2.101 AND 3.104.

4.1.3. Submission Timelines

This BAA shall remain open for three (3) years from the date of publication on www.fbo.gov. Although the Government may select proposals for award at any time during this period, it is anticipated that the funding for this program will be committed only during the first and second selection periods as stipulated on the first page of this Proposer Information Pamphlet (PIP). Proposers may submit a full proposal in accordance with the instruction provided herein at any time up to the proposal due date for the respective phase.

All submitted proposals will be reviewed. In order to be considered during the first round of funding, for Phase 1, full proposals must be submitted to DARPA, 3701 North Fairfax Drive, Arlington, VA 22203-1714 (Attn.: **BAA 06-29**) **on or before 12:00 Noon Eastern Daylight Saving Time, October 27, 2006.**

The proposal window and instructions for the second round of selections, for Phase 2, will be announced through a revision of this BAA that will occur during the latter half of Phase 1 execution.

Proposals submitted under this BAA may be either mailed or hand-delivered.

Mailing address: DARPA
ATTN: **BAA 06-29**
3701 North Fairfax Drive
Arlington, VA 22203-1714

For hand deliveries, the courier shall deliver the package to the DARPA Visitor Control Center at the address specified above. The outer package, as well as the cover page of the proposal, must be marked "**CORONET BAA 06-29**."

4.2. PROPOSAL CONTENTS AND FORMATTING CHARACTERISTICS

All submissions must be in the following format—nonconforming proposals may be rejected without further review. Proposals must be on single-sided pages, written in English, and with 1-inch margins (left, right, top, and bottom) in each page. A page is defined as being no larger than 8.5" by 11.0". (Accordion-style foldouts will be counted as multiple pages equivalent to the expanded size.) Paper copies of proposals should be stapled or submitted in loose-leaf binder, not bound. The Technical Proposal shall contain no smaller than 12 point font type. The Cost Proposal shall contain no smaller than 8 point font type. Larger font type for the Cost Proposal, up to 12 point font type, is desired, where appropriate.

4.2.1. Phase 1 Proposal Format

Proposals shall consist of two volumes. Volume I, CORONET Phase 1 Technical and Management Proposal, and Volume II, CORONET Phase 1 Cost Proposal. Volume I consists of 4 mandatory sections (I-IV) and one optional section (V). Page limitations for each section are identified in the section description below in braces {}. The page limitation for proposals includes all figures, tables (except the table of contents), and charts. All pages that exceed the maximum page limit specified may be removed and not be reviewed or considered in evaluation. Volume I, excluding Section V, is limited to 64 pages. There are no page limitations for Volume II and Section V of Volume I.

Your technical and cost proposals should conform to the guidance provided in Paragraphs 1.3. (Program Scope and Funding) and 1.4. (Period of Performance for Scope) of this BAA. Proposers should refer to those sections for information on how to scope and segment their technical and costs proposals.

Proposers must submit:

- one (1) original of the full proposal and
- six (6) copies of the full proposal and
- two (2) electronic copies of the full proposal
    - Electronic copies must be on a CD-ROM.
    - Each disk must be clearly labeled with **BAA 06-29**, proposer organization, and proposal title (short title recommended)
    - Electronic copies of the proposal must be in MS-Word readable application. Cost proposal spreadsheets should be submitted in an MS Excel-readable format.
        - Exceptions: the three relevant papers included in Volume I, Section V may be in .pdf format. No other items may be submitted in .pdf format.

4.2.1.1. <u>Volume I, Technical and Management Proposal</u>

Section I. Administrative {limited to 3 pages}

1. {1} Cover sheet to include:
    a. BAA number (**BAA 06-29**)
    b. Lead Organization Submitting proposal
    c. Type of business, selected among the following categories: "LARGE BUSINESS," "SMALL DISADVANTAGED BUSINESS," "OTHER SMALL BUSINESS," "HBCU," "MI," "OTHER EDUCATIONAL," or "OTHER NONPROFIT"
    d. Contractor's reference number (if any)
    e. Other team members (if applicable) and type of business for each
    f. Proposal title
    g. Technical point of contact to include: salutation, last name, first name, street address, city, state, zip code, telephone, fax (if available), electronic mail (if available)
    h. Administrative point of contact to include: salutation, last name, first name, street address, city, state, zip code, telephone, fax (if available), electronic mail (if available)
    i. Funds requested from DARPA for the Base Effort (Phase 1) and the amount of cost share (if any)
    j. Date proposal was prepared
2. {1} Official transmittal letter

3. Table of Contents. The Table of Contents should be keyed to the page numbers of the proposal sections {not included in page count}.
4. {1} A one slide summary of the proposal in PowerPoint that quickly and succinctly indicates the main objective, key innovations, expected impact, and other unique aspects of the proposal.

Section II. Detailed Proposal Information {limited to 47 pages}:
This section provides the detailed discussion of the proposed work necessary to enable an in-depth review of the specific technical and managerial issues. Specific attention must be given to addressing both risk and payoff of the proposed work that make it desirable to DARPA.

1. <u>{1} Executive Summary of the proposal</u>: This section should succinctly describe the uniqueness and benefits of the proposed approach relative to the current state-of-art and alternate approaches. Define the problem/challenge that this innovative claim will address and the effort's technical goals. Explain how this proposal addresses this problem differently than current approaches and the significant gains due to its uniqueness.
2. <u>{3} Innovative claims for the proposed research.</u> This section is the centerpiece of the proposal. It should succinctly describe the uniqueness and benefits of the proposed approach relative to current state-of-the-art and alternate approaches.
3. <u>{3} Deliverables</u> associated with the proposed research and the plans and capability to accomplish technology transition and commercialization will clearly address how the proposed effort will meet the goals of the program. Include in this section all proprietary claims to results, prototypes, intellectual property, or systems supporting and/or necessary for the use of the research, results, and/or prototype. (SEE SECTION 3.4, INTELLECTUAL PROPERTY.) If there are no proprietary claims, this should be stated. NOTE: For purposes of completing section 3.4, Intellectual Property, this information will not be counted in the proposal's page count.
4. <u>{3} Statement of Work (SOW)</u> written in plain English, outlining the scope of the effort and citing specific tasks to be performed and specific contractor requirements.
5. <u>{3} Cost, schedule and milestones</u> for the proposed research, including estimates of cost for each task in each year of the effort, for Phase 1, and total cost and company cost share, if applicable.
6. <u>{21} Detailed technical rationale, technical approach, and constructive plan</u> for accomplishment of technical goals in support of innovative claims and deliverable production. Includes a thorough quantitative discussion of relevant technical information and a detailed plan. This section should clearly explain: What you are proposing; why you are proposing this approach; why you believe it can be done now; and the importance or effect if successful. This should include:

   o High-level description of the network node architecture, including the hardware architecture of the optical network elements

   o High-level description of the architecture of the control and management planes, including the fault management, configuration management and performance monitoring functions as well as security concerns

   o High-level description of the proposed algorithms and protocols for meeting the goals of the CORONET program. This should include a basis for confidence, whether derived via analysis or simulations, that the proposed system will be sufficient in meeting the metrics specified in Section 2.5. This should also include proposed metrics related to performance and economics (in addition to those of Section 2.5) for evaluating the algorithms/protocols/architecture, as appropriate.

7. {5} Comparison with other ongoing research indicating advantages and disadvantages of the proposed effort
8. {3} Discussion of proposer's previous accomplishments and work in this or closely related research areas
9. {2} Description of the facilities that would be used for the proposed effort. If conducted with operational forces, what agreements/coordination has been made or will be required to meet this requirement
10. {3} Formal teaming agreements that are required to execute this program and a brief synopsis of all key personnel. A clearly defined organization chart for the program team that includes, as applicable, the:
    a. programmatic relationship of team members;
    b. unique capabilities of team members;
    c. task responsibilities of team members;
    d. teaming strategy among the team members; and
    e. key personnel along with the amount of effort to be expended by each person during each year.

Section III. Network Study {limited to 7 pages}
The proposal includes a network study, in order to gauge preliminary algorithmic performance with respect to network resiliency against multiple failures. While proposers need to present a global network topology as part of the proposal, note that for purposes of the proposal, the study is restricted to the CONUS portion of the network.

The following assumptions should be made in this network study:

- The study should be limited to Design Scenario 1 (see Table 3)
- The network study should be restricted to the CONUS portion of the topology; i.e., the 20 Tb/s of demand in Scenario 1 should be assumed to be CONUS-only traffic
- The traffic should be assumed to be 100% fixed (i.e., no churn). The proposal must include a discussion on how the results will be affected by traffic churn.
- The network services are as specified in Table 2
- All of the specifications and limits in Table 6 hold, except the restoration times in this table are not relevant for this network study
- The metrics of Table 5 and Table 7 are not a part of this network study

Discussion on the network design and the results should be included in this section of the proposal. At a minimum, this should include:

- Illustration of a global network topology that meets the requirements of Table 1, as well as a representation of the topology (i.e., cities, links, link distances) in MS Excel-readable format

- A traffic distribution corresponding to the network study for Demand Scenario 1, for S_o and P_o; the demand endpoints, service type, data rate, and protection level should be provided in MS Excel-readable format. (Note: the distribution should not be included as part of the text of the proposal.)
- The following six plots (see Section 2.5.2):
  - Spare capacity ratio vs. the parameter *P*, with the parameter *S* fixed at S_o
  - Spare capacity ratio vs. the parameter *S*, with the parameter *P* fixed at P_o
  - *W* vs. the parameter *P*, with the parameter *S* fixed at S_o
  - *W* vs. the parameter *S*, with the parameter *P* fixed at P_o
  - The number of links with extra capacity vs. the parameter *P*, with the parameter *S* fixed at S_o
  - The number of links with extra capacity vs. the parameter *S*, with the parameter *P* fixed at P_o
- For the nominal network design (S_o and P_o), statistics on:
  - average (weighted by bit rate) and maximum end-to-end working path length for all demands
  - average and maximum link utilization in terms of wavelengths
  - average wavelength fill rate (for the wavelengths carrying IP services)
  - average wavelength bypass of the IP routers across the network (i.e., the percentage of wavelengths that enter nodes from the network side that do not enter the IP routers)
  - average wavelength bypass of nodes across the network (i.e., the percentage of wavelengths that enter nodes from the network side that directly pass through without being dropped or regenerated (either OEO or OOO))
  - average size of an IP router across the core network, and maximum size of an IP router
  - any other statistics relevant to the particular proposed solution that can shed light on the economics and the technological feasibility of the solution

Section IV. Estimate of Phase 2 Scope, Cost and Schedule {limited to 7 pages}:

This section should include a high-level proposal for Phase 2. The proposal should be clear, concise and convey a understanding of the technical and programmatic challenges that must be addressed in Phase 2. At a minimum the description should include such things as possible teaming arrangements, program scope, high-level development and test plans, a high-level transition plan, estimated period of performance and ROM cost (to be used by DARPA for budgeting purposes only).

Section V. Additional Information (Optional) {no page limit}:

This section may include a brief bibliography of relevant technical papers and research notes (published and unpublished) that document the technical ideas upon which the proposal is based. Copies of not more than three (3) relevant papers can be included in the submission. The submission of other supporting materials along with the proposal is strongly discouraged and will not be considered for review.

4.2.1.2. Volume II, Phase 1 Cost Proposal – {No page limit}

1. A cover sheet to include:
   a. Name and address of Proposer (include zip code);
   b. Name, title, and telephone number of Proposer's point of contact;
   c. Award instrument requested: cost-plus-fixed-fee (CPFF), cost-contract--no fee, cost sharing contract--no fee, or other type of procurement contract (specify), agreement, or other award instrument;
   d. Place(s) and period(s) of performance;
   e. Funds requested from DARPA for the Base Effort, each option (if any) and the total proposed cost; and the amount of cost share (if any);
   f. Name, mailing address, telephone number and Point of Contact of the Proposer's cognizant government administration office (i.e., Office of Naval Research/Defense Contract Management Agency (DCMA)) (if known);
   g. Name, mailing address, telephone number, and Point of Contact of the Proposer's cognizant Defense Contract Audit Agency (DCAA) audit office (if known);
   h. Any Forward Pricing Rate Agreement, other such Approved Rate Information, or such other documentation that may assist in expediting negotiations (if available);
   i. Contractor and Government Entity (CAGE) Code,
   j. Dun and Bradstreet (DUN) Number;
   k. North American Industrial Classification System (NAICS) Number [NOTE: This was formerly the Standard Industrial Classification (SIC) Number]; and,
   l. Taxpayer Identification Number (TIN).
   m. All subcontractor proposal backup documentation to include items a. through l. above, as is applicable and available).

2. Detailed cost breakdown to include:

   a. Total program cost broken down by month and government fiscal year (GFY) [Note: Government Fiscal Year runs from October 1st to September 30th] and further broken down by major cost items as follows:
      i. Direct Labor – Individual labor category or person, with associated labor hours and unburdened direct labor rates;
      ii. Indirect Costs – Fringe Benefits, Overhead, General and Administrative Expense, Cost of Money, etc. (Must show base amount and rate);
      iii. Travel – Number of trips, number of days per trip, departure and arrival destinations, number of people, etc.
      iv. Subcontract – A cost proposal as detailed as the Proposer's cost proposal will be required to be submitted by the subcontractor. The subcontractor's cost proposal can be provided in a sealed envelope with the Proposer's cost proposal or will be requested from the subcontractor at a later date;
      v. Consultant – Provide consultant agreement or other document which verifies the proposed loaded daily/hourly rate;
      vi. Materials – Should be specifically itemized with costs or estimated costs. An explanation of any estimating factors, including their derivation and

application, shall be provided. Please include a brief description of the Proposer's procurement method to be used;

vii. Other Direct Costs – Should be itemized with costs or estimated costs. Backup documentation should be submitted to support proposed costs.

b. Costs of major program tasks and major cost items by year and month;

c. An itemization of major subcontracts (labor, travel, materials and other direct costs) and equipment purchases;

d. A summary of projected funding requirements by month; and

e. The source, nature, and amount of any industry cost sharing, if applicable. Where the proposer offers options for Phase 1, any cost sharing of the options should be identified separately as part of the cost estimate for each option.

3. Supporting cost and pricing information in sufficient detail to substantiate the summary cost estimates above. Include a description of the method used to estimate costs and supporting documentation. Provide the basis of estimate for all proposed labor rates, indirect costs, overhead costs, other direct costs and materials, as applicable.

4.2.2. Phase 2 Proposal Format

The specific requirements and instructions for the format of the Phase 2 proposal will be defined through an update of this BAA, anticipated to occur in the latter half of Phase 1 execution.

5. PROPOSAL EVALUATION

The criteria to be used to evaluate and select proposals for this program are described in the following paragraphs. Each proposal will be evaluated on the merit and relevance of the specific proposal as it relates to the program goals. The evaluation criteria in descending order of importance are: (a) Innovative Technical Approach; (b) Network Study In The Proposal; (c) High-Level Phase 2 Proposal; (d) Work Plan, Management Approach and Past Experience; (e) Potential Contribution and Relevance to the DARPA Mission; and (f) Cost Reasonableness and Realism. In accordance with FAR 35.016(e) the primary basis for selecting proposals for award shall be technical, importance to agency programs, and funds availability. Cost reasonableness and realism shall also be considered to the extent appropriate as described herein. Proposals may be evaluated as they are received, or they may be collected and periodically reviewed. The following are descriptions of how the Government will evaluate these factors.

5.1. INNOVATIVE TECHNICAL APPROACH

The objective of this criterion is to evaluate: the depth of problem understanding demonstrated in the proposal, the efficacy of the overall network architecture, the novelty of algorithms and protocols, the quantitative basis of confidence for the proposed work, and the proposal's realism.

The degree of problem understanding must be evident in the proposal and its sensitivity to the central issues identified in the BAA. Indicators of the depth of understanding will be sought, such as an awareness of the full scope of the program goals and consideration of related work, historical dead-ends, and risks that the proposer plans to mitigate. Other indicators include the extent to which the proposal milestones are appropriate to the task, and the nature of the milestone review criteria that the proposal identifies. Proposals that do not reflect awareness of relevant technical literature will be viewed unfavorably.

The proposed network architecture will be evaluated to determine if it meets the program vision as articulated in this PIP. Factors characterizing a suitable architecture include innovation, generality, scope, scalability (in cost, footprint, and power consumption), and compatibility with existing and emerging standards.

Proposed algorithms and protocols will be evaluated with attention to such issues as scalability, robustness, stability and convergence, consideration of network control and management overhead, and other relevant factors.

The proposal's quantitative basis of confidence in the approach and its benefit will be evaluated. The proposal should identify the analytic and simulation results that serve as a foundation for the plan. The proposal should articulate a logical and quantitative argument to the effect that the result will achieve the stated program metrics.

Many aspects of the proposal's realism will be evaluated. This includes, but is not limited to: the proposed network node architecture, including the architecture of the optical network elements, and the technical and economic assumptions used to justify the overall network architecture.

5.2. NETWORK STUDY IN THE PROPOSAL

The objective of this criterion is to evaluate preliminary algorithmic performance with respect to network resiliency against multiple failures, and the level of expertise demonstrated in the network study.

The efficiency of the network restoration scheme will be evaluated in terms of the spare capacity ratio defined in Section 2.5.2; in particular, how does the attained ratio compare to the 75% metric given in Table 6 for the nominal network design (S_o and P_o). Also, the scheme will be evaluated with respect to its sensitivity to the parameters S and P, and to the number of links that require an extra fiber pair.

The level of expertise will be evaluated in terms of modeling capabilities, and the realism of the global network topology and the traffic distribution included in the proposal.

5.3. HIGH-LEVEL PHASE 2 PROPOSAL

The objective of this criterion is to evaluate: the adherence to the goals articulated in this PIP; the appropriateness of the high-level development and test plans; the credibility of the proposed commercial transition strategy, and whether the proposed solution addresses realistic challenges of network deployment; and reasonableness and realism of the estimated effort and ROM cost.

5.4. WORK PLAN, MANAGEMENT APPROACH AND PAST EXPERIENCE

The objective of this evaluation criterion is to assess whether the statement of work and the management approach is likely to lead to a successful outcome. Factors to be evaluated include viability of the proposed work plan, roles and responsibilities, experience and multidisciplinary expertise, and opportunities for team interactions enabled by the management plan.

The overall research agenda and timeline, including specific intermediate criteria and milestones, should clearly relate to theoretical obstacles that must be overcome. Proposers should specify objective criteria for assessing progress of the team effort in terms of deliverables and intermediate objectives.

Project roles and responsibilities must be clearly defined. The roles of the prime and other participants required must be clearly distinguished and pre-coordination with all participants (including Government facilities, if any) fully documented. Management plans must demonstrate superior Government visibility into and interaction with key technical activities and personnel, and single point of responsibility for contract performances. Key personnel in the proposal will be included in the contract with clauses for contract termination in event of unapproved key personnel changes.

The team must have the experience and multidisciplinary expertise necessary to address the technical challenges. The proposed team must have sufficient experience to address the technical challenges, develop research breakthroughs, and manage cost and schedule. Similar efforts completed/ongoing by the proposer in this area should be fully described including identification of other Government sponsors, if any.

Opportunities for team interactions should be enabled by the management plan in order to obtain the full benefit of the multidisciplinary skill set.

5.5. POTENTIAL CONTRIBUTION AND RELEVANCE TO THE DARPA MISSION

The objective of this criterion is to establish a strong link between this work and the DOD mission. Factors to be evaluated include potential contribution to the national technology base, transition potential, intellectual property provisions, documentation, and dual use opportunities.

The potential contributions of the proposed effort with relevance to the national technology base will be evaluated. Specifically, DARPA’s mission is to maintain the technological superiority of the U.S. government and prevent technological surprise from harming our national security by sponsoring revolutionary, high-payoff research that bridges the gap between fundamental discoveries and their military use.

It is not necessary that the proposed work be immediately usable in military systems; however, a traceable path to transition opportunities should be identified. As mentioned throughout this PIP, the possibility of the transition of the results of the CORONET program to commercial carriers is an important goal of Phase 2 of the program. Proposals that identify a promising forward path for potential deployment of a revolutionary technology, despite the influence of legacy constraints, will be viewed favorably.

Intellectual Property restrictions on the objective system may be viewed unfavorably to the extent that affects the Government's ability to transition the program.

Plans for the documentation and publication of system documentation and source code will be carefully scrutinized, since documentation is essential to evaluation and future transition.

5.6. COST REASONABLENESS AND REALISM

The objective of this criterion is to establish that the proposed costs are reasonable and realistic for the technical and management approach offered. A secondary goal is to assess the proposer’s practical understanding of the effort. This will be principally measured by cost per labor-hour and number of labor hours proposed. Factors to be evaluated include the financial plan's completeness, transparency, modularity, necessity, and sufficiency:

- Completeness concerns whether the proposed costs reflect the program activities described in the PIP and BAA.
- Transparency concerns the degree to which costs can be mapped to specific technical tasks and milestones.
- Modularity refers to whether or not the separable components are clearly identified and can be evaluated in the case that partial award is deemed appropriate.
- Necessity refers to whether or not costs are required by the proposed tasks.
- Sufficiency refers to whether or not the costs described can be accomplished within the stated budget.

The evaluation criterion recognizes that undue emphasis on cost may motivate proposers to offer low-risk ideas with minimum uncertainty and to staff the effort with junior personnel in order to be in a more competitive posture. DARPA discourages such cost strategies. Cost reduction approaches that will be received favorably include innovative management concepts that maximize direct funding for technology and limit diversion of funds into overhead.

Awards will be made based on the overall best value to the Government. DARPA reserves the right to make awards to some, all, or none of the proposals received. Additionally, DARPA reserves the right to accept proposals in their entirety or to select only portions of proposals for award, and to award without discussions. In the event that DARPA desires to award only portions of a proposal, negotiations will be opened with that offeror.

NOTE: PROPOSERS ARE CAUTIONED THAT EVALUATION SCORES MAY BE LOWERED AND/OR PROPOSALS REJECTED IF SUBMITTAL INSTRUCTIONS ARE NOT FOLLOWED.

ATTACHMENT A: Sample R&D and Financial Status Report

(1) R&D STATUS REPORT
This brief narrative, not to exceed five pages in length, shall contain the following:

(i) For first report only; the date work actually started.
(ii) Description of progress during the reporting period, supported by reasons for any change in approach reported previously.
(iii) Planned activities and milestones for the next reporting period.
(iv) Description of any major items of experimental or special equipment purchased or constructed during the reporting period.
(v) Notification of any changes in key personnel associated with the contract during the reporting period.
(vi) Summary of substantive information derived from noteworthy trips, meetings, and special conferences held in connection with the contract during the reporting period.
(vii) Summary of all problems or areas of concern.
(viii) Related accomplishments since last report.
(ix) Fiscal status, to include reporting of summary level financial data in the following format: (next page)

R&D STATUS REPORT
PROGRAM FINANCIAL STATUS

Work Breakdown		Cumulative to Date			At Completion
Structure or Task Element Remarks	Planned Expend	Actual Expend	% Budget Compl	Latest At Compl	Revised Compl Estimate

Subtotal: ____

Management
Reserve: ____
Or
Unallocated
Resources: ____

TOTAL:
==

Note: Budget at completion changes only with the amount of any scope changes. (Not affected by underrun or overrun)

Based on currently authorized work:

Is current funding sufficient for the current fiscal year (FY)? (Explain in narrative if "NO")

YES NO

What is the next FY funding requirement at current anticipated levels?

$ ____________________

Have you included in the report narrative any explanation of the above data and are they cross-referenced?

YES NO

(2) FINAL REPORT

This report shall document the results of the complete effort and should be delivered at the completion of the contract. If the Government chooses to exercise the options under this contract, the due date for the final report is extended accordingly. Title pages shall include a disclaimer worded substantially as follows:

> "The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied, of the Defense Advanced Research Projects Agency or the U.S. Government."

The Final Technical Report summary shall include:

Task Objectives
Technical Problems
General Methodology (i.e., literature review, laboratory experiments, surveys, etc.)
Technical Results
Important Findings and Conclusions
Special Comments
Implications for Further Research
Standard Form 298, September 1988

(b) Reports delivered by the Contractor in the performance of the contract shall be considered "Technical Data" as defined in Section I contract clauses entitled "Rights in Technical Data – Noncommercial Items" and "Rights in Noncommercial Computer Software and Noncommercial Computer Software Documentation."

(c) Bulky Reports shall be mailed by other than first-class mail unless the urgency of submission requires use of first-class mail. In this situation, one copy shall be mailed first-class and the remaining copies forwarded by less than first-class.

(d) All papers and articles published as a result of DARPA sponsored research shall include a statement reflecting the sponsorship. In addition, a bibliography of the titles and authors of all such papers are to be included in the Final Technical Report

(1) The cover or title page of each of the above reports or publications prepared will have the following citation:

Sponsored by
Defense Advanced Research Projects Agency
STO
Program:
ARPA Order No. XXX, Program Code: XXX
Issued by DARPA/CMO under Contract No.:

(2) The title page shall include a disclaimer worded substantially as follows:

> "The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied, of the Defense Advanced Research Projects Agency or the U.S. Government."

(3) All technical reports must (i) be prepared in accordance with American National Standards Institute (ANSI) Standard Z39.18; (ii) include a Standard Form 298, August 1998; and (iii) be marked with an appropriate Distribution Statement.