



Department of Justice, Federal Bureau of Investigation
Property Procurement and Management Section
935 Pennsylvania, N.W., Washington, DC 20535

RFI TECHNICAL REFRESH PROGRAM (TRP) COMPUTERS AND O&M SERVICES

Sources Sought Notice

In response to: FBI-06-TRP-01

Dated: April 25, 2006

**Submitted
May 23, 2006**

**Prepared by: Nortel Government Solutions Incorporated
Tax ID # 54-1339972 – D&B #15-177-0955**





TABLE OF CONTENTS

1	COMMENTS.....	2
1.1	INDUSTRY IT STANDARDS.....	2
1.2	TRAINING.....	2
1.3	COOP PLANNING.....	2
1.4	CLIENT/SERVER ARCHITECTURE.....	3
1.5	STANDARDS BASED TECHNOLOGY.....	3
2	QUESTION & ANSWER.....	4
2.1	QUESTION 1.....	4
2.2	QUESTION 2.....	4
2.3	QUESTION 3.....	5
2.4	QUESTION 4.....	6
2.5	QUESTION 5.....	6
3	ALTERNATE SOLUTIONS.....	7
3.1	SPLIT MULTI-LINK TRUNKING.....	7
3.2	10GB ETHERNET.....	7
3.3	SIP COLLABORATION.....	8
3.4	INTERNET PROTOCOL VERSION 6.....	9

1 COMMENTS

Overall, the Draft Technology Refresh Project Functional Requirements Document is a well developed, detailed, and complete roadmap of the functionality the Federal Bureau of Investigation (FBI) would like their infrastructure to support over the next four years. By partnering with the final RFP selected vendor(s), the FBI will implement a comprehensive four-year transformation. With proper planning, the FBI can also evolve the requirements into the model for continuous change, which will enable infrastructure to grow and progress well beyond the initial contract period.

With some reviews of technology and FBI policies, and by working with industry, the FBI could also bring about some technology changes that could benefit the Bureau and other mission partners within the Department of Justice and Intelligence Community. Our comments are focused on five areas, and they are as follows:

1.1 INDUSTRY IT STANDARDS

Implementation of components of the Information Technology Infrastructure Library (ITIL) would improve overall management of your IT infrastructure. Additionally, management of the IT infrastructure will also improve by applying the Carnegie Mellon University's Capability Maturity Model Integration (CMMI). Both of these management and process models are complementary and are among the highest of industry standards for IT management. Companies and government agencies that have followed the methodologies of these models have improved their service and reliability of their IT infrastructure, and also reduced their long-term costs. By ensuring that your Project Managers for the eventual RFP are knowledgeable in either Project Management Body of Knowledge (PMBOK) or certified as Project Management Professionals, you can assure that the project proceeds in an orderly manner, and will meet well-defined milestones till the completion of the project.

1.2 TRAINING

The Enterprise Operations Center (EOC) needs to ensure that users do not go into "Future Shock" or suffer from "Information Overload" by introducing changes into the infrastructure at a pace that overwhelms the user's abilities to compensate. Just a change of Windows Office packages can cause a great deal of confusion and distress among workers. To minimize these risks, a plan for training users on the new emerging technologies needs to be part of the TRP. The better employees are prepared for the changes to the infrastructure the fewer general type trouble tickets EOC will face. By providing training, you enable employees to embrace change and see it as an opportunity, not a roadblock. Time and money spent on training is not wasted, but it can take a while for its benefits to be seen.

1.3 COOP PLANNING

The COOP planning that has been done by the FBI shows the detail and effort that has gone into the design of the back-up strategy. The only concern is that there is no mention of the method of off-site data storage. Due to the large volume of data that is stored on the backup systems even at a Field Office it can take several hours to restore a system from tape. It is advisable to perform monthly critical data back-ups over the network to provide a fail-safe to the onsite data storage.

1.4 CLIENT/SERVER ARCHITECTURE

The idea presented in the document that thinner client/server architecture requires less IT support staff is not necessarily true. The fact that a thinner client has no local configuration is somewhat true. However, depending on how the user is to interface with the server side of the system; it could require a great deal of customization in providing the user with the same look and feel that would be native in a standard desktop PC environment. Trouble shooting issues can be just as complex with thin client as with thick. The larger the numbers of thin client users the greater the demand will be on the infrastructure's communication backbone. Also, the dependency on server hardware and hardware fail over is increased. A very long look has to be paid to both arguments of thick client versus thin client because the net sum gain of one over the other could be zero.

1.5 STANDARDS BASED TECHNOLOGY

By reviewing available and emerging technologies; possible revision of FBI policies; and working alongside industry the FBI could foster some technology changes that can benefit the Bureau and other justice and intelligence agencies. We would encourage the FBI to work with industry to utilize standards-based technology and best practices specifically in the areas of SIP, IPv6, 10GB, and SMLT.

2 QUESTION & ANSWER

2.1 QUESTION 1

Q: Is the FBI's view of technical refresh viable and is it keeping pace with industry?

A: As to “is Technical Refresh Project (TRP) viable”, yes, but only if a very strong and well-defined Release Management process to support the change to the infrastructure accompanies it. Release Management is a function of the Information Technology Infrastructure Library (ITIL) model. By ensuring that change is managed, it will become easier and more cost effective to roll out new technologies throughout the time span covered by the TRP and well beyond.

As to “is the TRP keeping pace with industry?”, over-all, no. If looking at just today’s technology then the answer would be yes, but looking at what will be available within the year, no. An analysis needs to be performed to determine what applications, COTS and government specific, are currently in use; are they the best solutions for how they are used; and who is using the software. Once that is completed a review of software needs to be performed to determine which applications have upgrades available and if the upgrade of these applications still fulfills the software’s initial requirements

Also, a review of policies may be in order, due to the fact that as of now IP Telephony is not allowed as a solution for voice communication.

2.2 QUESTION 2

Q: Are there additional services and alternative product configurations that are currently deployed by Industry and not features in the Draft document?

A: Yes, there are several areas:

KVM - The draft document mentions the concern over the use of KVM switches and the need to remain with the PS/2 format for keyboards and mice. This is no longer a limitation since there are several vendors offering USB-based KVM switches that even support the sharing of USB devices like printers and scanners. By reviewing these devices it might be possible in some areas to reduce costs by having the ability to share at the desktop an inkjet printer between the secure and non-secure systems.

IP Telephony technology - There are several companies offering very robust and secure IP-based telephone systems that utilize dedicated or converged network backbones; interface voice mail to email; and provide collaboration capabilities including secure instant messaging and desktop video conferencing. If there is way to work with industry to develop a system that meets the security concerns of the FBI for IP Telephony, then this should be pursued. By integrating voice mail and voice-over-IP technology into the FBI infrastructure, it can become an additional investigative tool available to field/special agents.

2.3 QUESTION 3

Q: Are there better ways to technically refresh the FBI's IT infrastructure?

A: Yes, by following the principles of Information Technology Infrastructure Library (ITIL).

Change Management - If the FBI does not have an established process for handling change to the IT infrastructure on a continual basis then the process needs to be developed before TRP begins. Over a four-year period it is likely that technologies will change, software will be revised, and hardware will be phased out by the manufactures. By having a process in place to handle these inevitable events, time and money can be saved. It will also be necessary to incorporate end-user business group's requests for new or revised functionality. Changes in focus and/or mission of some or all the business groups can play a role in the type of changes that may need to be incorporated into TRP over a four-year span.

Configuration Management - If the FBI does not have an established process for managing the infrastructure's configuration once TRP begins then the process needs to also be developed before it does. One method to reduce the cost of Configuration Management is to reduce the types of hardware in the environment. Currently the FBI has at a minimum 15 configurations of workstations (3 security levels and 5 hardware models). If at the end of the four-year period this can be reduced to 2 hardware models per each security level (6 configurations), the cost savings should be apparent. By maintaining the configurations over the four-year period, the FBI will also be able to improve break-fix and trouble tickets because less time will be spent on issue resolution. Applying the same configuration controls to laptops, printers, and communication equipment, reduces the over-all complexity of the infrastructure from a management standpoint. In the realm of software, limiting the types and versions allows for easier file sharing, template design, and trouble-shooting. For example: Listed under software are 6 versions of the Java Runtime environment, the newest version is SE v1.4.2_09. This version supports all of the previous 5 versions listed. There should only be a need for the newest version. By setting defined standards and communicating those to developers, this process would improve development time, lower integration and testing costs, and improve adoption of new applications by the end-users.

Asset Management - If the FBI does not have an established process for handling asset recovery and redeployment then before TRP begins the process needs to be developed. By recovering most of the older platforms from the first year of TRP and redeploying the workstations to sites that would not get the newer equipment of TRP, FBI can still raise the over all level of service to the IT infrastructure. This will also increase the longevity of the infrastructures hardware assets. A surplus supply can be established to cover break-fix issues and/or emergency response concerns.

Release Management - If the FBI does not have an established process for performing package releases to the infrastructure, then before TRP proceeds, the process needs to be established. Once TRP has begun, the need for continual infrastructure transformation becomes dependant on the ability to effectively handle changes that need to be made due to problem resolutions generated by Enterprise Management; functional requests from the end-user business groups; and technology changes forced by manufacturer life cycles. Release Management works to ensure you have strong Configuration Management and Change Management, and works along side Asset Management to ensure that any new assets are introduced into the infrastructure correctly and with the minimum of difficulty.

At the end of TRP the organization should have a well developed and in-place Life Cycle management process that covers all types of change to the infrastructure.

2.4 QUESTION 4

Q: Industry, you have seen the IT resources currently in use at the FBI, can you identify better solutions?

A: Yes, there have been and will be by end-of year several major upgrades to software listed in the Draft document.

Windows 2003 - The Windows 2003 OS, including Exchange 2003, is a proven stable upgrade to Windows 2000 OS and offers enough features to warrants its inclusion it a technology refresh.

Windows Vista - Microsoft will be releasing Windows Vista, the replacement to Windows XP, to remain current the development of this platform should be included into the scope of the TRP.

Microsoft Office 2003 - Office 2003 is a much more robust version of the Microsoft Office suite, and Outlook 2003 works best with Exchange 2003.

WordPerfect - Corel has released a new version of WordPerfect it is WordPerfect Office X3.

Adobe Reader - Adobe has release Version 7.0.7 of their reader.

Pointsec - The newest version of Pointsec is version 6.0, released in December of 2005. This version gives a much wider range of hardware choices of smart card readers.

McAfee Virus Protection - McAfee now offers an out of the box Enterprise solution to virus protection that covers both servers and workstations, as well as assists with enterprise management.

2.5 QUESTION 5

Q: Industry, you have seen the IT resources currently in use at the FBI, can you identify weaknesses or performance issues?

A: Section 3 of this response describes in detail four standards-based technologies that could provide improved performance to the IT resources.

SMLT - Split Multi-link Trunking provides an additional level of protection against failures for the data network core. SMLT enables node redundancy by allowing MLT links of link-aggregated groups to be dual-homed across a pair of aggregating devices, which are then individually connected to edge network devices. Without changing the edge topology or devices, SMLT enables core redundancy for mission critical applications including voice and video.

10 Gigabit Ethernet - The size and amount of information has grown to fill the capability of existing Gigabit networks. The industry has responded by developing the next generation of network backbone speed at 10Gbs. Several manufacturers are in the process of implementing the 10Gb technology for SAN and other transport applications.

SIP – Provides integrated, always on communication and collaboration capabilities for true convergence.

IPv6 – Lays the foundation for transparency, security and support for future services.

3 ALTERNATE SOLUTIONS

3.1 SPLIT MULTI-LINK TRUNKING

Split Multilink Trunking (SMLT) can provide redundancy and full utilization of existing riser connections between a core Ethernet or routed Ethernet switching environment without requiring a refresh of the existing wiring closet configuration. SMLT delivers support for the 802.1ad standards-based Trunking that typically exists in wiring closets in 2, or 4 x 1Gb links to each core switch (typically 2) in an ‘all ports active’ mode. Whereas deployments using Spanning Tree, or a technology other than SMLT, make only half the ports active to a single core switch, effectively cutting the capacity of these networks in half, and reserve the remaining half for redundancy (with considerable Spanning Tree convergence times). SMLT utilizes all ports during normal operation and half of the available ports in the event of a failure. Leveraging technology similar to what has been developed for the next generation 10Gb Ethernet networks, SMLT provides sub-second failure rates with the LAN/Campus environment. This is critical for Voice over Internet Protocol (VoIP) and other delay/jitter sensitive applications.

For instance, typical resilient Ethernet networks consist of wiring closet (edge) switches dual homed to network center aggregation (core) switches in a building or campus. This implementation requires the use of the Spanning Tree Protocol to protect the network against loops. While the Spanning Tree Protocol in any form (IEEE 802.1D/w) suits this purpose, it comes with certain limitations. Although it does protect against loops, it is not optimized so that networks can fully utilize all links (blocked ports) and at the same time be redundant. As a result, link aggregation technologies have become popular to improve link bandwidth and/or protecting against link failures.

SMLT leverages IEEE 802.3ad, currently part of IEEE 802.3-2002 clause 43, standardized link aggregation protocol to improve link utilization while at the same time providing protection against link failures. It improves the level of Layer 2 resiliency with nodal protection extensions, link failure protection and flexible bandwidth scaling. SMLT achieves this by allowing edge switches using 802.3ad to dual home to two SMLT aggregation switches. The SMLT is transparent to these attached devices.

SMLT inherently avoids loops because of its superior enhanced-link aggregation-control-protocol, and eliminates the use of the IEEE 802.1D/w Spanning Tree protocols. Two aggregation switches appear as a single device to edge switches that are dual homed to the aggregation switches. The aggregation switches are interconnected using an InterSwitch Trunk (IST), over which they exchange addressing and state information, permitting rapid fault detection and forwarding path modification. Both links are fully active in normal operation and, during a failure; the traffic on a failed link switches to the remaining link in less than one second.

For more information on SMLT, please refer to the following IETF draft:

<http://www.ietf.org/internet-drafts/draft-lapuh-network-smlt-06.txt>

3.2 10GB ETHERNET

LAN PHY and WAN PHY interconnections for core networks are generally accepted in the enterprise market as the logical next step to increase bandwidth on a Gigabit-based core backbone. Several manufacturers have begun offering high-bandwidth, low-latency, and low-cost 10Gb backbone switches to connect aggregation switches in the core.

We recommend evaluating 10Gb as part of a distributed-routing environment. Distributed-routing environment improves the ability to increase core bandwidth and provide higher redundancy levels without adding complexity to the network. In a distributed routing and switching environment, wiring closets can connect to multiple core switches and utilize all the bandwidth on the uplinks/risers to the core without the limitations of a centralized routing core. In effect, by adding high density, low-cost stackable Ethernet routing switches across the entire backbone, the network core provides each IP flow a shortest path first option from any access port in the network.

By interconnecting each routing switch with a network 'backplane' fiber that extends a routing switch's resources up to 80 kilometers (standard optics dependant) apart. A single 'master' switch manages all the resources in the backbone, regardless of location and shares the MAC and routing tables with each other core Ethernet switch in the network. By extending a single layer 2 and layer 3 core across the enterprise, the backbone becomes a single hop transport. This drastically reduces routing table calculations and decisions for routing packets. It also localizes the VLAN to the network core, not just to the first core switch that VLAN traffic enters. In a chassis based implementation, one core switch transport can introduce between 5 microseconds of delay up to 1200 microseconds of delay between core switches; transport delay notwithstanding. The 10Gb switches provide 4.58 microsecond Gigabit-Gigabit transport and can maintain this low latency over considerable distances (i.e., less than 10 microseconds for 10 km).

With a 35 to 55% cost advantage over existing campus core deployments, the cost reductions achieved by deploying a 10G distributed routing environment can be used to fund the applications that new backbone would easily able to accommodate. We can provide more information upon request

3.3 SIP COLLABORATION

The Technology

Session Initiation Protocol (SIP) forms the basis of the industry-standard IP-centric converged communications architecture by serving as a signaling mechanism in establishing a wide variety of sessions. In this context, a session is any interactive communication that takes place between two or more entities over an IP network, from a simple two-way telephone call or an instant message exchange, to a collaborative multimedia conferencing session. The Internet Engineering Task Forces rfc2543 defines SIP as follows:

SIP is an application-layer control (signaling) protocol for creating, modifying and terminating sessions with one or more participants. These sessions include Internet multimedia conferences, Internet telephone calls and multimedia distribution. Members in a session can communicate via multicast or via a mesh of unicast relations, or a combination of these.

SIP invitations used to create sessions carry session descriptions that allow participants to agree on a set of compatible media types. SIP supports user mobility by proxying and redirecting requests to the user's current location. SIP is not tied to any particular conference control protocol; it is designed to be independent of the lower-layer transport protocol and can be extended with additional capabilities.

SIP exhibits the following attributes: Ease of programming via text-based message formats; Re-use of HTTP for syntax in message headers and cause codes; Flexible addressing via e-mail like addresses; Expendability by leveraging Internet building blocks (e.g., DNS, LDAP, and RADIUS) and by allowing application servers (e.g., conferencing); Scalability via highly-distributed networking; and Security by leveraging IP security protocols (e.g., SSL, IPSec) and functionality (e.g., firewalls).

In so doing, SIP eliminates the pain points users, knowledge workers, and other face with the need to communicate and/or collaborate including: Managing multiple contact numbers and inboxes (the lack of service ubiquity); Losing productivity when away from the office (the lack of geographic flexibility); and

Using disparate systems (e.g., telephones, room video conferencing, e-mail, file servers) to communicate across teams.

Benefits

Integrated communications are made up of asynchronous communications (e-mail, voicemail, short message services) and synchronous communications (IM, voice, video, and application sharing) combined with presence and location intelligence. Integrated implies a seamless user experience across all these media; always-on implies a level of reliability not generally associated with platform supporting applications such as e-mail, workflow and document handling.

Ultimately, a SIP name will become the primary address that needs to appear on a business card. SIP-based multimedia collaboration and virtual office capabilities will eliminate geographic boundaries. Always-on integrated communications will provide a unified interface to a full range of multi-media communication capabilities. Secure Instant Messaging (including archiving) will compliment multi-media conferencing capabilities by providing real-time status signaling.

SIP will accomplish this by supporting the following applications:

SIP-based collaboration: Multimedia, video and audio conferencing, secure instant messaging, IM chat, Web collaboration, application sharing, file exchange, Web push and co-browsing, white boarding and clipboard sharing.

SIP-based mobility: People-centric session set-up, presence-enabled personalization and customization, network-enabled find-me follow-me and location based services, adaptive capabilities to user device and network connectivity.

SIP-based productivity and information interactions: Network-based directory, network-wide session screening, and management including click-to-call handling and session logs, automatic and user-ser presence management, real-time information exchange tagging of all session with subject fields and context sensitive session rejection—all made available to all devices.

Reduced Costs

SIP applications provide both strategic payback and tactical ROI by increasing productivity and reducing costs across the enterprise infrastructure. For example: SIP immediately reduces the costs that are typically associated with conferencing and collaboration by decreasing the dependence on disparate systems and service providers to deliver these capabilities. Furthermore, it improves productivity by making these services available to users at any point on the network at any time. A user can be assigned a dedicated conference bridge and from this initiate multimedia conferences, including the exchange of large amounts of data and video.

3.4 INTERNET PROTOCOL VERSION 6

Internet Protocol Version 6 (IPv6) will support the deployment of Internet services to the vastly growing number of devices likely to join the network in the years ahead. Begun in 1989, the effort to develop IPv6 was triggered by the forecasted exhaustion of the available IPv4 address space. It has grown over the years to include many new features for the IP architecture. This includes providing flexibility for new applications and support for the expected vast number of new users. IPv6 is believed to be essential to the deployment of secure, robust networking for large-scale military and civilian applications. Limited commercial IPv6 services are currently available in some international backbones and to end users in research and education networks worldwide. IPv6 is designed to be a complete replacement for the network layer of the IP protocol stack. This layer incorporates the addressing for the packets of information and handles the routing of these packets.

IPv6 delivers: 128-bit addresses as compared to the 32-bit addresses of IPv4 in use today. Global connectivity without NATs, Integration of IPsec security capabilities; Additional support for deployment of services requiring QoS; Simplified Administration; and the flexibility needed to accommodate future services.

IPv6 is designed to relieve many of the constraints of IPv4 and will underpin the continuing expansion of the Internet. The larger address space of IPv6 is particularly important to the very large number of new mobile devices, some enabled through the SIP technology (described above), and that wire line and wireless broadband access will bring to the Internet over the next few years. The widespread adoption of Network Address Translation (NAT) technology as a stopgap solution for address space exhaustion has already affected the ability to send a packet from anywhere-to-anywhere across a large network without modification. IPv6 will aid in restoring this “transparency,” providing much better support for person-to-person, person-to-machine, and machine-to-machine communications. IPv6 also facilitates end-to-end security with the ubiquitous implementation of IP security in end points.

IPv6 serves to reestablish the any-to-any connectivity and automated recovery from even major multi-point network failures. NAT translates the IP addresses of packets as they pass across the boundary between privately addressed networks and public networks such as the Internet, allowing non-unique, private IPv4 addresses to be reused in multiple domains. With NAT, packets no longer travel across the network unmodified. This removes the transparency of the network, limiting the applicability of IP Security (IPsec) and impeding the development of person-to-person communications. Deployment of IPv6 will address these limitations and reduce the management burden imposed by the use of NAT, especially where IPv4 addresses are very scarce. IPv6 conformance in hosts and routers requires support for IPsec. Taken together with transparency, it is designed to allow a converged network based on IPv6 to deliver secure services end-to-end on whatever path a packet is routed.

The deployment of IPv6 is being carefully managed to avoid the problems with IPv4 routing table size that have arisen over the last few years. The hierarchical allocation of IPv6 addresses, mirroring the packet transport hierarchy, helps to reduce the number of routes needed at each level in the hierarchy, but rules out some of the IPv4 traffic engineering techniques used to improve robustness by multi-homing sites on the network.

We are promoting a structured-resilient architecture for IPv6 networks. This provides the full benefits of IPv6 distributed routing under major failure conditions while leveraging the advantages of well-established resilient technology, e.g., SMLT, to support fast recovery in day-to-day operations. It also helps avoid the challenges of deploying IPv6, including: making a smooth transition from IPv4 through coexistence to an eventual all-IPv6 network, providing rapid network recovery under a broad range of potential failure conditions; and balancing the need for connectivity with the imperative of security.

Compared with IPv4, IPv6 incorporates many features that will make it easier to roll out the emerging applications that are fueling demand for addresses. IPv6 packets can be carried over MPLS Label Switched Paths, Layer 2 and Layer 3 VLANs, and Virtual Private Networks. They support classification, policing, shaping, and routing of IPv6 packets across these technologies. We view IPv6 as a key enabler of a future robust, converged network capable of providing security, connectivity, and failure recovery, balancing normal operations and responsiveness to major events. We propose implementing IPv6 with Layer 2 switches that will provide IPv6 traffic classification and prioritization, Layer 3 devices that will provide IPv6 routing, transition and coexistence capabilities, and Layer 4-7 devices that will provide classification, filtering and load balancing for IPv6 traffic.