

**Envisioning the California Executive Branch's
Next Generation Enterprise Network**

By

**J. Clark Kelso
Chief Information Officer
State of California
June 12, 2007**



STATE CHIEF INFORMATION OFFICER

J. Clark Kelso (ckelso@pacific.edu)
3455 Fifth Avenue
Sacramento, CA 95817
(916) 739-7302 / (916) 739-7072 (fax)

June 12, 2007

During my five-year tenure as State CIO, I have periodically issued papers setting forth my vision for particular aspects of the Executive Branch's information technology program. *See, e.g., Government Services on the Web: "California In-Touch"* (May 10, 2006); *Re-Alignment of Responsibility for the Management of the State's IT Resources and Infrastructure* (May 12, 2003); *Information Technology Procurement, Management and Operations – Preliminary Report* (July 1, 2002). These papers have provided transformational visions, forming the basis for subsequent planning and substantial changes in State IT.

Today's paper deals with another key transformation that is already underway: our transformation to the next generation of enterprise networking, a transformation that recognizes and takes full advantage of the convergence of voice, data and video communications, supports the burgeoning use of wireless technologies, and builds upon the consolidation of our two general-purpose data centers into the Department of Technology Services (DTS).

We are well positioned to take advantage of converging and developing network and telecommunications technologies by virtue of the Calnet II contract, administered by DTS, which gives all agencies convenient access to the most modern telecommunications and network services at competitively-procured rates.

However, Calnet II is only a procurement vehicle. It is the responsibility of the State's IT leaders to advise agency executives on the best use of the Calnet II contracts. This document is one piece of that puzzle. In addition to this document, the IT Council will soon be considering a draft *California Statewide Telecommunications Plan – Pathway to a Connected California*, a document under development since November 2006. That strategic plan will provide additional guidance. Finally, DTS is undertaking a detailed examination of its own network infrastructure and will be releasing a plan later this year for its improvement.

Taken together, these documents and their combined recommendations should assist all agencies in making the best decisions possible about the acquisition and management of telecommunications and network technologies.

Clark Kelso
Chief Information Officer
State of California

I. Executive Summary

Introduction

The California Executive Branch has been working steadily over the last three years to plan for and begin executing an across-the-board refresh of its information and telecommunications technologies. We have successfully concluded a leveraged procurement for modern telecommunications services. We have successfully undertaken strategic sourcing for common information technology commodities. We have begun IT modernization projects in several of our largest departments. Other departments are conducting modernization assessments of their IT programs. We have come to a consensus that we need to expand our use of Executive Branch-wide applications – true enterprise applications to serve the common business needs of Executive Branch agencies.

Against the backdrop of these improvements, it is now time for us to modernize our network architecture and infrastructure. This will not be an easy or quick journey. Our existing telecommunications and network architecture and infrastructure have been built over a twenty or thirty year period, and built largely without any Executive Branch-wide or enterprise vision. At a time when all agency budgets are stretched, when cyber-security threats continue to rise, and when substantial enterprise applications are on the horizon immediately before us, we no longer can afford the ad hoc, agency-centric solutions of the past. An enterprise network approach is an imperative.

“Enterprise Network” Defined

For our purposes, an “enterprise network” is defined as follows:

An enterprise network consists of all of the hardware, software, people, policies, and practices that an organization employs to electronically transmit information from one person or machine to another person or machine consistent with a uniform set of organizational standards for security, availability, reliability, speed and cost-effectiveness.

There are two important features of this definition. First, an “enterprise network” encompasses more than just the hardware and wiring that constitutes the physical infrastructure. It includes the software, people, policies and practices that bring network services to life. Second, the logical scope of an organization’s network is essentially set by the organization’s standards for security, availability, reliability, speed and cost-effectiveness. Uniformity of those standards across the enterprise defines the scope of the network.

The Executive Branch’s Enterprise Network: CSGnet II

In light of this definition, one of the foundational questions to consider is whether the Executive Branch should try to establish a single, general purpose enterprise network that spans all or nearly all agencies within the Branch (essentially recognizing the entire Executive Branch as an enterprise), or, in the alternative, whether each agency should establish its own agency-specific

enterprise network, thereby permitting enterprise networks to proliferate across the Executive Branch. At present, we have an ad hoc, mixed environment, some portions of which are managed centrally by the Department of Technology Services and other portions of which are managed by each agency.

Establishing for the Executive Branch a centrally or jointly managed, general purpose enterprise network (a “California State Government Network II” or “CSGnet II”) to provide secure network services and connectivity to virtually all agencies and programs is likely to be the best approach to improve cost-effectiveness and security and to ensure alignment between enterprise applications and network architecture. Although these considerations strongly point in the “CSGnet II” direction, we recognize that further review and validation of this conclusion is needed as we undertake our additional network planning since the full costs and implications of consolidating networks cannot be known absent consideration of the details of particular proposals and alternatives.

Additionally, before we make a final decision to migrate from our current, essentially ad hoc architecture to a CSGnet II, we need to establish branch-wide security policies applicable to all agency networks to ensure a baseline level of security. An enterprise network will never be properly secured absent branch-wide security policies. When that policy development has been completed, we should then examine whether a further consolidation of authority over enterprise networks, or creation of joint authority between DTS and its enterprise network customers, would substantially improve cost-effectiveness, security and alignment.

Under this approach, all agencies would be required to use CSGnet II for their network services. Exemptions from the mandatory use of CSGnet II should be granted only by the State CIO and only to those few agencies that can clearly and convincingly establish markedly different business requirements for security, availability, reliability, speed or cost-effectiveness for their networks.

Network Governance

Enterprise networks call out for new governance to organize our decision-making and policy-making processes. The type of governance we need over our enterprise networks encompasses much more than decisions only from a technology perspective. There will be fundamental business, policy and funding issues that will arise, particularly at the inception of CSGnet II, in addition to substantial technical questions about architecture, availability and security. Accordingly, the Technology Services Board should consider establishing a broadly inclusive “Enterprise Network Committee” to be charged with overall responsibility for developing policies regarding CSGnet II, the Executive Branch’s next generation enterprise network.

In addition to governance for policy development, we may need to establish joint governance over network operations and management encompassing DTS and its enterprise network customers. Joint governance over operations would ensure that all parts of the CSGnet II family are working in harmony and that important issues of enterprise security, as well as individual agency business needs, have an appropriate forum for discussion and resolution.

II. A Framework for Visioning Enterprise Networks

A. Introduction

Creating an overall vision for enterprise networks capable of handling the next generation of communications needs for California's Executive Branch is a daunting challenge. To begin, the size and scope of California's many government agencies, and the diversity of their business functions, presents a major conceptual hurdle. Given this scope and diversity, is a single, coherent vision even possible?

Moreover, our organizational diversity and complexity must now be mapped against recent developments in the computer and telecommunications fields that cloud attempts at clear visioning. Where once there were relatively clear distinctions between data and voice traffic, and between types of information processing devices and telecommunication devices, today we see a convergence of voice, data and video traffic over networks and a rapidly progressing march of information processing components into virtually all devices. Yesterday's telephone has become another type of computer. Our cell phones and blackberry's come with operating systems that make them increasingly acceptable as Internet browsers. What used to be a relatively dumb desktop printer is today an intelligent, networked device capable of multiple functions. Wireless connectivity supports our increasingly mobile workforce. With all of these devices connecting to each other, the business uses for networks have expanded along with the technologies. How do we design enterprise networks for these divergent uses?

Next, our conception of a state network must take account of the explosive growth and increasingly all-encompassing role of the Internet. Because of the Internet, nearly all networks are now connected together globally, and any attempt to impose an organizational vision upon networks could quickly be overwhelmed by the practical importance of Internet connectivity. How should we distinguish between "our" networks and the Internet's global network?

Finally, the pace of innovation in network, telecommunication and computing technologies shows no signs of slowing down, and with these technologies converging on each other, the pace of change is likely to accelerate. As noted in a very useful report prepared for the State CIO by the University of Southern California's Collaborative on E-Governance, "[n]ew and emerging technologies daily present new challenges along with important new opportunities for government, and the State has to stay on top of this technology wave or risk being drowned by it." *California's Telecommunications Strategy: Business and Management Challenges*, p. 2 (January 2007). How do we establish a sustainable vision when the technology environment is changing so rapidly?

With such fundamental questions facing us, it is best to revisit and revalidate some first principles that should guide our thinking. Those foundational issues include:

- What Is an "Enterprise Network"?
- Ideally, How Many Enterprise Networks Should the Executive Branch Maintain?
- Which Agencies, If Any, Should Be Served By Separate Enterprise Networks?
- How Will the Executive Branch Govern Its Enterprise Networks?

Answers to these questions will position the Executive Branch to adopt a more detailed strategic plan for telecommunications and network services. Drafts of such a plan are already in preparation, so the time for resolving foundational issues has arrived.

B. Defining the “Enterprise Network”

Amidst the technologically complex and changing environment, and in an organization as large and diverse as California’s Executive Branch, what does it mean to own, maintain and use an “enterprise network”? What is an “enterprise network”?

Giving a precise definition to “enterprise network” is not a simple task. In general, a definition should be a concise statement that explains the meaning of a word or phrase, or that specifies the essence of the thing being defined so that one can readily and reliably determine whether an object falls within the class of things being defined. A good definition avoids circularity, over- or under-breadth, and unnecessary ambiguity or obscurity.

If we define “enterprise network” too broadly, then nearly all of our information technology devices will be viewed as being connected to the same network. For example, we could consider the Internet to be one very large network encompassing everything that is connected to it. This definition is too broad. It does not permit us to distinguish between “our” network space, which we manage to serve our business needs, and “other’s” network spaces. At a bare minimum, we need a definition that will describe “our” network space, and distinguish our network space from “other’s” network spaces, so we can properly secure our data and communications. The Internet is not a secure network, and its architecture does not serve all of our needs. We need a narrower, more focused definition.

From a technologist’s perspective, a network is “an interconnection of three or more communicating entities” (Telecommunications: Glossary of Telecommunication Terms), and a computer network is a “network of data processing nodes that are interconnected for the purpose of data communication” (*Id.*). These technical definitions are both too broad and too narrow for our purposes – too broad because anything connected to anything else would constitute a network, and too narrow because the sole focus is on the technology. These definitions do not reflect the organizational context in which the issue of enterprise networks arises. An organization needs to know where its enterprise network begins and ends, what resources are deployed to manage its enterprise network, and how to distinguish between its enterprise network (where communication can, by and large, be trusted) and outside networks (where communication generally is untrusted).

Instead of defining “enterprise network” from a technology perspective, we need to define “enterprise network” from an organizational perspective that takes account of the organization’s business needs, including its security needs, and encompasses all of the resources that are necessary to make a network useful. The following definition is more suitable for our purposes:

An enterprise network consists of all of the hardware, software, people, policies, and practices that an organization employs to electronically transmit information

from one person or machine to another person or machine consistent with a uniform set of organizational standards for security, availability, reliability, speed and cost-effectiveness.

There are two important features of this definition. First, the definition encompasses more than just the hardware and wiring that constitutes the physical infrastructure. It includes the software, people, policies and practices that bring network services to life. While it might be possible in theory to separate the hardware and software in one organizational bucket, draw upon policies and practices from another organizational bucket, and have the network managed by people drawn from a third organizational bucket, this separation of functions and accountability is likely to lead to inefficiency and functional gaps in service. A well run network is one that combines all of the pieces together under one management team.

Second, although it may not be apparent at first reading, pursuant to this definition, the logical scope of an organization's network is essentially set by the organization's standards for security, availability, reliability, speed and cost-effectiveness. The uniformity of those standards across the enterprise defines the scope of the network. In this way, an organization's *business needs* for particular levels of security, availability, reliability, speed and cost-effectiveness are the primary basis for determining the boundaries or perimeters between networks. These standards define what is "ours" and may be trusted, and what is "not ours" and must not be trusted.

From a technology perspective, networks are also usually divided into categories based upon the scale or extent of reach of the network, giving us Local Area Networks (LANs), Campus Area Networks (CANs), Metropolitan Area Networks (MANs), and Wide Area Networks (WANs). In addition to these hard-wired networks, we also must anticipate substantially increased reliance on wireless networks to provide access to State systems (this increased reliance includes both state employees using non-state-owned wireless networks for remote and mobile access to state systems, a type of access that creates special security problems, as well as an increase in the number of state-operated wireless networks). The complete architecture of Executive Branch networks includes all of these network types. Nearly all agencies have some form of Local Area Network, usually managed by the agency itself, which then connects to a Wide Area Network, most commonly managed by DTS. The use of wireless is clearly on the rise everywhere, including open wireless networks, state-operated wireless networks, and wireless access to state systems through PDAs, Blackberry's, and other mobile devices.

C. How Many Enterprise Networks Should the Executive Branch Maintain?

1. The Problems With Network Proliferation

Because our definition of "enterprise network" encompasses all of the resources necessary to acquire and manage a network and defines the scope and perimeter of the enterprise network, one of the foundational questions to consider is whether the Executive Branch should try to establish a single, general purpose enterprise network that spans all or nearly all agencies within the Branch (essentially recognizing the entire Executive Branch as an enterprise), or, in the alternative, whether each agency should establish its own agency-specific enterprise network.

In theory, each agency could establish its own enterprise network with a zone of trust that extends no further than the boundaries of that agency. Inter-agency communications would be treated as communications between untrusted networks. Under this approach, enterprise networks would proliferate across the Executive Branch.

Are there any problems with permitting enterprise networks to proliferate across the Executive Branch? The answer is a definite yes, and for three inter-related reasons:

- Cost-effectiveness;
- Alignment between business needs, enterprise applications and network architecture; and,
- Security.

a. Cost-Effectiveness.

Permitting every department to build and manage its own enterprise network is likely to result in substantially increased costs without a clear business case justification or commensurate return on investment. Each department would be required to obtain technical resources to acquire and manage its own enterprise network, either by hiring employees or contracting for the necessary expertise or buying services from a provider. Each department would be required to build or buy its own network backbone infrastructure. Each department would be required to purchase firewalls and other network perimeter defenses to protect its network from all other networks (including all other state agency networks, since in this model, every network outside of a department would be treated as an untrusted network). Each department would be required to manage that perimeter, building in scores of pathways, tunnels and exceptions to security rules to permit legitimate state traffic into the departmental network. And in most cases, each of these expenditures will essentially duplicate similar expenditures by most other agencies. We will, in effect, be buying the same network over and over again, throughout the Executive Branch.

In summary, letting each department build its own enterprise network is likely to increase capital costs, costs for ongoing maintenance and operations of that capital investment, and security-related costs incurred because of the additional complexity inherent in this decentralized architecture.

In general, cost-effectiveness is best achieved by aggregating common IT and network infrastructure, where possible, so that economies of scale can be achieved in the acquisition and management of that infrastructure. My paper on *Re-Alignment of Responsibility for the Management of the State's IT Resources and Infrastructure* (May 12, 2003) made the basic case for consolidation of the State's two general purpose data centers. As a result of the great leadership shown by DTS's executive team, and the hard work of both DTS's staff and customers, consolidated-related savings have surpassed \$40 million annually. Consolidation works when carefully planned and implemented. The same fundamental set of economic and management arguments apply to network consolidation.

Admittedly, this assessment of cost-effectiveness requires further validation, since the full costs and implications of consolidating networks cannot be known absent consideration of the details of a particular proposal and comparison of its costs against the costs of alternative architecture.

As DTS moves forward with its network planning efforts, it must undertake this more granular analysis. Preliminarily, however, there are strong reasons to suspect that the total costs of operating a small number of enterprise networks (e.g., under 5) is likely to be substantially less than the total costs of operating scores of enterprise networks.

b. Alignment.

An enterprise network should align with an organization's business needs for security, availability, reliability and speed. Proper alignment between business needs and network characteristics improves cost-effectiveness. As a generality, if an organization employs a large number of complex applications that must operate across the entire enterprise, those applications will be best supported by a single network where all of the touch points between users and computing resources are within the zone of trust. Only by employing a single enterprise network will it be possible to optimize the network for the applications being used, the number of users, the amount of traffic on the network, the distribution of network traffic over time, and the level of security required. Trying to run multiple, complex enterprise applications over scores of separately managed networks, each of which has firewalls protecting from untrusted outsiders, will dramatically increase the costs of building and managing the networks, as well as the costs and complexity of managing the enterprise applications. Technical feasibility will be difficult to achieve under this approach.

The Executive Branch already promotes utilization by all departments of a small number of Executive Branch-wide enterprise applications, including CalStars (the State's accounting system) and CalAters (a travel claim system). In addition, many of the State's most important statewide applications benefit from a centrally-managed, statewide network, including DMV's vehicle registration, the child support system, EDD's statewide employment-related programs, the statewide automated welfare system, and the child welfare system / case management system, to name a few of the bigger applications. Finally, there are common network applications, such as Statewide email and Internet access. Internet access already constitutes close to 50% of the traffic on state networks, and that usage will only increase with time, especially considering what is now becoming the widespread transmission of video and audio information over the Internet.

As set forth in the *Statewide Information Technology Strategic Plan* (November 2006), the future holds out the prospect of a vastly increased number of Executive Branch-wide enterprise applications, including the 21st Century Project's payroll system, a statewide ERP system (Fi\$Cal) encompassing budgeting, accounting, procurement, and asset management, statewide portal functionality, and health IT systems, to name just a few of the major projects that will be coming online in the next 3 to 7 years. In essence, we are moving steadily towards enterprise-wide applications for nearly all basic administrative functions.

The Executive Branch needs to have enterprise network functionality and architecture that is aligned with and supports these enterprise applications. In light of these enterprise applications, no department can legitimately claim to be an island unto itself. In most cases, the commonalities that bind us together are greater than our differences. Accordingly, we should be limiting the

overall number of enterprise networks as much as we can which will help us better align our networks with our overall business needs.

c. Security.

Perhaps the most important consideration favoring development of a relatively small number of enterprise networks is management of the Executive Branch's overall cyber-security risks at reasonable cost.

With the growth of the Internet over the last decade, and now the dramatically increased use of wireless technologies, threats to the security of computer systems and networks have become ubiquitous. For identity thieves, intellectual property pirates, purveyors of pornographic and other illegal material, spammers and ordinary hackers out for a good time, large organizational networks have become a primary target. Securing a network against assaults is an extraordinarily complex task requiring the concentrated efforts of virtually everyone in an organization who touches a computer. Although hackers are routinely probing Executive Branch networks directly for unguarded points of entry, our existing enterprise networks have successfully defended against direct attacks. Unfortunately, there are as many indirect points of entry as there are users and devices attached to the network. Unless restrictions are imposed on use, users can navigate to dangerous spots on the Internet where malicious code can make its way onto a network, and use of wireless networks and devices presents additional challenges.

Even without an Internet connection, malware can be introduced into a network environment by users who can unwittingly introduce foreign elements into a network by bringing data or programs from home on CDs, DVDs or thumb drives. Once introduced into the network, such programs can reproduce themselves throughout the network.

These are not merely hypothetical musings. For example, early in 2007, an Executive Branch agency experienced a security incident involving a worm infection. At the time of the infection, the commercial cyber-security companies did not have a signature for the worm or a patch to protect against it, nor any instructions on how to eradicate it. The worm spread quickly through the department's network-connected devices. To limit the possibility of the worm skipping through the Executive Branch's enterprise network to other departments, the decision was taken to shut down the department's Internet connection and the pathway to other departmental networks. Through the extraordinary efforts of the department's staff, the worm was brought under control within 24 hours. However, the incident serves as a very stark reminder about the reality of our vulnerabilities.

Cyber-security risks increase in proportion to network complexity. An organization that builds and manages 20 networks that, for business reasons, must interoperate is an organization that has substantially greater cyber-security risks than an organization that builds and manages a single organizational network and imposes uniform security standards throughout the organization. In the organization with 20 networks, the overall network is only as secure as the weakest of the 20 networks. That means the organization has to invest heavily in security in all 20 networks, and the likelihood of a gap or failure among those 20 networks is statistically higher than the

likelihood of a gap or failure in only 1 or 2 heavily secured networks. In short, trying to manage 20 networks is a riskier venture than trying to manage a smaller number of networks.

2. CSGnet – The Executive Branch’s Primary Enterprise Network

Consistent with its overall mission, the Department of Technology Services manages a general-purpose enterprise network, known as CSGnet, for use by all departments. CSGnet is a statewide private network operated by DTS, using DTS owned and managed hardware, with circuits procured through the CALNET contracts. CSGnet provides agencies with access to the Executive Branch’s enterprise-wide applications, access to the Internet, and connectivity between state and local government entities for certain applications. DTS also provides departments with remote access services, including access through wireless networks. CSGnet is the Executive Branch’s primary enterprise network.

As explained above, the Executive Branch’s overall goal should be to reduce to a bare minimum the number of enterprise networks. In order to achieve this goal, agencies should be required to use CSGnet for anything other than Local or Campus Area Networks, and all Internet connections should be made through CSGnet. As discussed further below, exceptions to the mandatory use of CSGnet will be granted by the State CIO only upon an extraordinary showing of business need.

CSGnet itself needs to undergo certain transitions to serve the Executive Branch’s future telecommunications and network needs. We need CSGnet to transition to a “CSGnet II” which will provide agencies with the full range of modern telecommunications and networks services as well as the capacity to stay abreast of new network technologies, including the burgeoning field of mobile and wireless devices. The first two steps towards a CSGnet II are as follows: First, DTS needs to take full advantage of the services now available on Calnet II and manage those services for the benefit of DTS’s customers. Second, DTS and agencies need to take certain steps to improve overall network security.

a. Calnet II Network Services

The Calnet II set of contracts provides Executive Branch agencies with convenient access to a full range of modern telecommunications and network technologies. Under Calnet II, agencies will be able to purchase services from four different “Modules” encompassing the following types of services: (1) Voice, Data and Video Services; (2) Long Distance and network based services; (3) Internet Protocol (IP) Voice, Data and Video Services; and (4) Broadband Fixed Wireless Access Services. Services under the first two modules are provided by AT&T; services under the second two modules are provided by Verizon.

Most of the Executive Branch’s telecommunications and network needs can be satisfied through the Calnet II contracts. It appears that DTS and other agencies could begin the transition to a CSGnet II almost exclusively by reliance upon purchases of additional network services from Calnet II. Under this approach, the Executive Branch’s primary enterprise network would be a joint public-private venture, owned, maintained and managed by the Calnet II contractors and

DTS. CSGnet II would become the primary source of network services to all Executive Branch agencies.

Achieving the right mix of state-owned and managed infrastructure versus vendor-owned infrastructure with DTS oversight will be one of the critical management decisions for us to make as we transition to CSGnet II. In a number of agencies, the state has developed significant expertise in certain areas of network management, and that expertise enables us better to tailor, tune and secure our network infrastructure to meet our business needs. By the same token, generic or commodity-based network and telecommunication services may well be best acquired from the private sector where there exists both additional economies of scale and greater flexibility in adopting new technologies. In all events, the state will remain deeply involved in designing, managing and securing its enterprise networks.

To achieve a smooth transition to CSGnet II, several changes must be made. The first change is internal to DTS's organization. Presently, DTS maintains two separate divisions with overlapping authority for networks. The Calnet II program is managed by the Statewide Telecommunications and Network Division (STND). DTS's separate Engineering Division manages CSGnet and DTS's relationships with its agency customers acquiring network services. These two divisions must be merged under common leadership. With the convergence of telecommunications traffic and data traffic on a common network infrastructure, it no longer makes sense to separate telecommunications from network engineering. Combining STND with the network portion of the Engineering Division will improve DTS's ability to manage and oversee the enterprise network.

The Calnet II modules do not currently proscribe any particular levels of network security. Since DTS will rely upon the Calnet II vendors to provide more comprehensive services for the Executive Branch's enterprise network, DTS needs to work collaboratively with its customers and the Calnet II vendors to architect a secure, available, reliable and cost-effective network that meets the State's business and program needs. DTS will establish Service Level Agreements most particularly in the security area.

b. Improving Network Security

Today DTS and agencies treat CSGnet as a self-contained Wide Area Network that departments connect to for certain purposes. CSGnet's scope does not extend into the Local Area Network (LAN) space; however, because agencies have connected LANs to CSGnet, we have created multiple enterprise networks – CSGnet and each of the agency Local Area Networks connected all together as though they were actually one enterprise network. Unfortunately, in today's cyber-security environment, this type of equivocation in network architecture is unacceptable.

DTS does *not* manage agencies' local area networks; they are managed by the agencies. Yet these Local Area Networks traverse CSGnet which leaves everyone connected to CSGnet more vulnerable to security incidents since overall security on a network is only as good as the weakest link in the network. In other words, every agency connected to CSGnet must rely upon good security practices being followed by *every other agency* connected to CSGnet. In recent years, many agencies have understandably decided that this risk is simply too great for their

business needs, and have begun to install defensive firewalls at the connection point to CSGnet. This approach will become increasingly expensive and difficult to maintain as we expand the number of Executive Branch-wide enterprise applications.

We need to undertake a thorough assessment of the security environment and security practices of the Local Area Networks. We should consider whether security could be improved in some Local Area Networks by making DTS responsible – either solely or jointly with the affected agency – for managing and monitoring the Local Area Networks. Admittedly, this would be a new role for DTS and a substantial re-alignment of responsibility for network infrastructure, but it may be a necessary re-alignment in the long run to improve security.

Additionally, today DTS does not own or manage the edge routers (telecommunications equipment that is located at the client’s facilities) for many of the customers of the former Teale Data Center; therefore, DTS does not have direct accountability for maintenance of all routers at the edge of CSGnet and these routers are critical points of entry into the enterprise network. Historically, Teale did not insist upon its ownership of the edge routers and was more receptive to requests for customer ownership and management of that equipment. By contrast, the Health and Human Services Data Center (HHWDC) insisted upon ownership and management by the data center of edge equipment. Ten years ago, these divisions of accountability for network security between the Teale Data Center and HHWDC and their customers were arguably not a serious problem since the security risks from external sources were much smaller then. But today, with the growth of the Internet and wireless access, and with cyber-security incidents rising exponentially around the world, we can no longer afford the luxury of decentralized security management of our enterprise network.

As the Department of Technology Services works to consolidate the networks deployed by Teale and HHSDC, it must take ownership and responsibility for managing all routers on the edge of the network. DTS will treat networks from agencies that do not transfer ownership and maintenance of edge routers as untrusted networks. All agencies have a responsibility to each other and to the public to cooperate with DTS in this change to support our enterprise network.

Before embarking upon such major re-alignment of network responsibilities, we must first develop uniform, statewide network security policies for all agencies to follow for all network environments. Once we have developed the needed policies, we will be in a better position to evaluate the pros and cons of involving DTS in managing and monitoring Local Area Networks. We know improved security is a significant advantage, but agency management of the Local Area Network also has advantages, particularly when trying to diagnose local failures resulting from the interaction of networks with local applications and conditions.

D. Which Agencies Should Be Served By Separate Enterprise Networks?

The considerations set forth above lead to the conclusion that the Executive Branch should establish a centrally or jointly managed, general purpose enterprise network – a CSGnet II – to provide secure network services and connectivity to virtually all agencies. Ultimately, all resources on that enterprise network should be managed by the Department of Technology

Services, with the possibility of joint responsibility with agencies for monitoring and maintenance of Local Area Networks.

Exemptions from the mandatory use of CSGnet or CSGnet II should be granted only by the State CIO and only to those few agencies that can clearly and convincingly establish markedly different business requirements for security, availability, reliability, speed or cost-effectiveness for their networks. The Department of Justice and the Department of Water Resources are two examples of agencies that have clearly and convincingly established such markedly different business requirements. A brief review of their separate enterprise networks should help establish a baseline for when exemptions may be appropriate.

1. Department of Justice Network

The law enforcement and litigation activities of the California Department of Justice (DOJ) are supported by large repositories of highly sensitive and confidential data regarding offenders and criminal activity as well as volumes of legal documents, communications and depositions. DOJ maintains and manages a large and complex network connecting local, regional, state and federal law enforcement agencies to facilitate the exchange of accurate, timely, and complete criminal justice intelligence, identification and information, using state-of-the-art computer technology. The Hawkins Data Center hosts the data for the California Law Enforcement Telecommunications System (CLETS) and other applications such as CalGang and CalPhoto, as well as the largest single online database of fingerprints.

DOJ's network of over 650 routers connects 6 Attorney General Offices, 10 Remote Bureau of Forensic Sciences laboratories, and 30 DOJ regional, firearms, gambling and task force offices. More than 600 local law enforcement agencies and a variety of interstate and federal agencies access the network throughout the day.

Federal and State mandates, as well as operational considerations, require that this information be limited to a "need to know basis" and that the accuracy of the information be protected from alteration. The information requirements of both law enforcement and litigation entities necessitate that specific subsets of traffic be prioritized to ensure timely access to critical data. This access can mean the difference between success and failure, apprehension and escape, or even life and death.

To meet these stringent security needs, DOJ has deployed a private network backbone over which it utilizes Multi-Protocol Label Switching (MPLS). The private nature of the network allows DOJ to prioritize its traffic to meet its business needs and maintain strict control on the access to its information resources. The physical separation of the network enables DOJ to effectively limit the connectivity between its network and other state networks. All access to public web sites and the Internet are carefully isolated and secured behind multiple firewalls.

DOJ's business needs (many of which are dictated by federal requirements) call for a unique set of standards for security, availability, reliability, speed and cost-effectiveness that essentially define the perimeter and characteristics of DOJ's network.

2. Department of Water Resources Network

Another example of an agency which needs a separate enterprise network is found in the Department of Water Resources. DWR's unique business in the water and electrical markets requires specific security requirements for network boundaries while also requiring special trusted business-to-business connections. All outside entities are viewed as "untrusted," and only those entities that have a specific and approved business requirement, and have been sufficiently secured, are allowed connection on a specified machine to machine basis.

To support its unique business requirements, DWR maintains the Supervisory Control And Data Acquisition Wide Area Network (SCADAWAN) that is used to monitor and control the State Water Project. The SCADAWAN is a wide area network connecting the numerous pumping plants, generating plants, Area Control Centers and the Project Operations Center located through out the State. Due to the critical nature of the SCADAWAN, this network has its security boundaries configured following many of the best practices for SCADA security. Many of these practices are now being mandated by the North American Electric Reliability Council Cyber Security Standards for Critical Cyber Assets associated with the Bulk Electric Systems.

The SCADAWAN is a closed network. Retrieval of business and operations information from the SCADA network is always through a bastion host or Information Gateway located on one of two Demilitarized Zones (DMZs). Access through the bastion host and its accessible data is also controlled on a user by user basis. Access from the business network is never allowed directly into the SCADA system. Redundant business to business encrypted and firewalled connections with the California Independent System Operator (CAISO) utilize a custom version of a bastion host which is called a Remote Intelligent Gateway (RIG). These connections, for electrical regulation control and meter data settlements, are necessary for DWR to participate in the deregulated electrical market. Access is only granted between uniquely specified host systems. All other traffic, ports and protocols are blocked.

DWR's unique business responsibilities and needs, particularly for security, justify maintenance of a separate DWR enterprise network.

E. Governance

Establishing and maintaining a CSGnet II along the lines set forth in this document can be successful only if there is broad-based, inclusive network governance to ensure that decisions about network architecture, performance, usage policies, funding and security are made in the best interests of all. CSGnet II cannot be built by DTS alone and simply offered as a service to agencies. Instead, we must from the beginning ensure that CSGnet II is "owned" and governed by the Executive Branch as a whole.

Over the last several years, we have taken several important steps to organize our decision-making around information technology issues. There are now three major IT governance bodies in the Executive Branch with distinctly different enterprise roles and responsibilities:

- The Information Technology Council, composed primarily of agency CIOs, advises the State CIO on overall IT planning and policy, primarily from a technology perspective;
- The Technology Services Board, composed of representatives from all Cabinet agencies and the State Controller, governs the Department of Technology Services and sets policy on enterprise services provided by the Department of Technology Services; and,
- The Enterprise Leadership Council, with a membership similar to but slightly broader than the Technology Services Board, provides a forum for Executive Branch agencies to discuss and resolve business issues related to enterprise-wide IT from a business perspective.

The type of governance we need over our enterprise networks encompasses much more than decisions only from a technology perspective. There will be fundamental business, policy and funding issues that will arise, particularly at the inception of CSGnet II, in addition to substantial technical questions about architecture, availability and security. Accordingly, the Technology Services Board – which sets policy over DTS’s enterprise services – should consider establishing a broadly inclusive “Enterprise Network Committee” to be charged with overall responsibility for developing policies regarding the Executive Branch’s next generation enterprise networks.

In addition to governance for policy development, we may need to establish joint governance over network operations and management encompassing DTS and its enterprise network customers. Joint governance over operations would ensure that all parts of the CSGnet II family are working in harmony and that important issues of enterprise security, as well as individual agency business needs, have an appropriate forum for discussion and resolution.

III. Conclusion

California's Executive Branch has been working steadily over the last three years to plan for and begin executing an across-the-board refresh of its information and telecommunications technologies. We have successfully concluded a leveraged procurement for modern telecommunications services, we have successfully undertaken strategic sourcing for common information technology commodities, we have begun IT modernization projects in several of our largest departments, and other departments are conducting assessments of their IT programs with a view towards modernization. We have also come to a consensus within the Executive Branch on the need for a substantial expansion in the utilization of Executive Branch-wide applications – true enterprise applications to serve the common business needs of Executive Branch agencies.

Against the backdrop of these improvements, it is now time for us to modernize our network architecture and infrastructure. This will not be an easy or quick journey, because our existing telecommunications and network architecture and infrastructure have been built over a twenty or thirty year period, and built largely without any Executive Branch-wide or enterprise vision. At a time when all agency budgets are stretched, cyber-security threats continue to rise, and substantial enterprise applications are on the horizon immediately before us, we no longer can afford the ad hoc, agency-centric approach of the past. An enterprise approach to our enterprise networks is an imperative.