

#### **4.2.7 Content Delivery Network Services (CDNS) (L34.1.4.6, M.6, C.2.4.6-2.4.6.1.3)**

***Qwest's CDNS combines our converged Internet Protocol-based network and our team member [REDACTED]'s content delivery platform to enable fast and reliable delivery of Web-based content globally.***

Qwest's Content Delivery Network Services (CDNS) provides a suite of capabilities that off-loads origin servers and delivers content on their behalf. Our CDNS combines the unparalleled Qwest Internet Protocol (IP) network capability with [REDACTED]. Qwest's CDNS extends worldwide through Qwest's international Internet presence and [REDACTED] extensive global infrastructure. Qwest and [REDACTED] have a proven record of working together to provide CDNS for Government and commercial clients, including award-winning service for the Internal Revenue Service (IRS).

Qwest understands and complies with the designated standards, connectivity requirements, and technical capabilities for CDNS. Both team members actively participate in a number of standards-related organizations, and have played an active role in bringing new standards to the market. For example, [REDACTED], developed Edge Side Include (ESI) standards, which have been widely recognized and adopted by vendors to provide dynamic edge-based content delivery and processing.

##### **4.2.7.1 Reserved (L.34.1.4.6 (a))**

##### **4.2.7.2 Reserved (L.34.1.4.6 (b))**

##### **4.2.7.3 Satisfaction of CDNS Requirements (L.34.1.4.6(c))**

The following three sub-sections describe how Qwest satisfies all of the capabilities, features, and interfaces for CDNS.

#### **4.2.7.3.1 Satisfaction of CDNS Capabilities Requirements (L.34.1.4.6(c), C.2.4.6.1.4)**

Qwest fully complies with all mandatory stipulated and narrative capabilities requirements for CDNS. The following text provides the technical description required per L.34.1.4.6(c) and does not limit or caveat Qwest's compliance in any way. Qwest fully supports the CDNS capabilities for Networx. The requirements are organized into content distribution and site monitoring/ server performance measurements..

##### **Content Distribution**

Our specific service offerings include static content download service, real-time streaming, and on-demand streaming. The approach to each is briefly described in the following paragraphs.

**Static Content Download Service:** For delivery of static site content, each end-user request is directed to an edge server via intelligent Domain Name Service (DNS). Upon receiving a request for content, an edge server retrieves the appropriate content HyperText Markup Language (HTML) page, image, document download, Secure Socket Layer (SSL) object, and Video on Demand file from a local cache and delivers the resulting content to the requesting user. If requested content is not in cache at the edge, it is retrieved from within the network or the origin site.



**RealTime Streaming:** For live streaming, depending on the format of the encoded media, encoders push the content—or the entry point pulls the encoded stream—into the CDNS. If the media being provided to Qwest is in a raw (un-encoded) state, the stream will need to be encoded in at least one of a variety of formats including, but not limited to, RealNetworks Real Media, Microsoft Windows Media and Apple Quicktime. The format chosen depends on the requirements of the Agency. Because encoding requirements (CODEC, stream rate, and acquisition method) may vary greatly from event

to event, this service will be custom designed. After the encoded stream is acquired, the reflectors on the CDNS network then route the stream from the entry point to the edge servers, while maintaining reliability and quality. Routing via the network uses sophisticated packet recovery techniques, including re-transmit and multiple path transmission, to guarantee the highest quality stream delivery to edge regions. The edge regions then distribute the streams to end-users. Pre-bursting significantly reduces the time necessary to buffer and start the stream.

**On-Demand Streaming:** For on-demand streaming, when a user clicks on a stream, they are routed to the optimal server. Encoding for on-demand streaming is handled in the same way as for real-time streaming except that the content is stored on CDNS-resident storage and typically not acquired “live”. Our CDNS employs a pull architecture: content is replicated in streaming server caches in response to user requests, as follows:

- The streaming servers cache only the content that the users view, not the entire file.
- The platform load balances among our streaming servers to ensure that no single machine is responsible for the delivery of all content.
- The HyperText Transfer Protocol (HTTP) protocol is employed to deliver on-demand streams to edge streaming servers.

### **Site Monitoring/ Server Performance Measurements**

Available through the Qwest Control Network Portal,   
 is a dashboard that provides a comprehensive collection of network management tools. This provides Agencies with instant visibility into their traffic, content, applications, and users on the Internet enterprise management systems, including HP OpenView and IBM Tivoli NetView. Agencies can integrate real-time



Qwest's CDNS provides a flexible failover service that provides multiple options for Agencies. Web sites that rely on centralized infrastructure often find that ensuring uptime is a continuous challenge. A typical solution involves mirroring a website at an alternate location; however, this approach creates additional capital and management costs. [REDACTED] site failover frees Agencies from the unnecessary capital and management costs associated with creating a failover solution, and is offered as a managed service.

**4.2.7.3.3 Satisfaction of CDNS Interface Requirements (L.34.1.4.6(c);**

**C.2.4.6.3)**

CDNS is an application layer service supported by the connectionless data services available with the Internet Protocol (IP) suite of protocols via the User-to-Network Interfaces (UNIs) discussed in the IPS Section of this technical volume. The CDNS provides data transfer from an origin server to the CDNS servers via IP. The service is available to all Agency servers reachable by IP.

**4.2.7.4 CDNS Quality of Service (L.34.1.4.6(d), C.2.4.6.4-C.2.4.6.4.1)**

Qwest will meet or exceed all Quality of Service (QoS) requirements for Networkx CDNS. **Figure 4.2.7-2** provides Qwest's Acceptable Quality Levels (AQLs).

**Figure 4.2.7-2 Qwest's CDNS Meets All Networkx AQLs**

Key Performance Indicator (KPI)	Service Level	Performance Standard (Threshold)	Acceptable Quality Level (AQL)	[REDACTED]
Availability (CDNS Network)	Routine	100%	100%	[REDACTED]
Latency (Static Content Download)	Routine	Mean = 1.5 sec	Mean < 1.5 sec	[REDACTED]
Grade of Service (Time to Refresh Content)	Routine	5 minutes	≤ 5 minutes	[REDACTED]
Time to Restore (TTR)	Without Dispatch	4 hours	≤ 4 hours	[REDACTED]
	With Dispatch	8 hours	≤ 8 hours	[REDACTED]

Qwest understands each of the Key Performance Indicators (KPIs) and meets and exceeds the requirements for availability, latency, Grade of Service (time to refresh content), and Time to Restore (TTR).

**Availability:** Qwest's CDNS leverages Qwest's highly reliable IP/MPLS network and [REDACTED]. [REDACTED] applies load balancing and mapping software at multiple levels. At the server level, [REDACTED] provides server redundancy throughout its infrastructure. At an enterprise level, [REDACTED] proprietary mapping algorithms continually monitor Internet status to determine the fastest, most reliable routing paths.

**Latency for Static Content Download:** The combination of Qwest's dedicated IP connectivity and [REDACTED] high performance platform with route optimization, flexible caching, and delivery from the optimal edge server ensures that we will meet and exceed the Networx latency performance requirements.

**Grade of Service (GoS) (Time to Refresh Content):** The Qwest Team's CDNS meets the requirement. While we can establish time to refresh content in five minutes or less, our CDNS provides much greater and flexible refresh levels.

**Time to Restore:** Our CDNS has been designed from the ground up to require minimal human intervention. It is a self-healing, autonomous network that facilitates our support of restoration requirements.

#### **4.2.7.5 Qwest's CDNS Exceeds Service Requirements (L.34.1.4.6(e))**

Qwest's CDNS exceeds both the failover service and redirection and distribution service requirements, and offers several additional capabilities to Agencies.

**Failover Service:** The Qwest Team exceeds this requirement through our ability to provide failover as a completely automated managed service,

with no hardware or software requirements, and to provide three failover options, based on the needs of the Agency. Failover options are detailed below.

**Failover Option 1—Failover to Qwest's [REDACTED]:** If an Agency wants to ensure that a complete origin site will be available to end users regardless of the health of the origin site and/or Internet connectivity, Qwest can establish a backup site [REDACTED]


[REDACTED]


[REDACTED] Qwest is able to ensure a much higher degree of reliability, security and performance, in addition to offloading the need for additional infrastructure. The Agency can also store a default page [REDACTED]

[REDACTED]

**Failover Option 2—Failover to Alternate Data Center (Mirror Failover):** In the event that an Agency wants to be protected against the failure of an origin site and is running a backup or alternate site, site failover can be directed to use the backup site in case the primary is unavailable. Note that the backup site may, at the discretion of the Agency, be different from the origin site—for example, a site containing reduced functionality or content. Upon receiving a request for a piece of content that must be obtained from the primary site and determining that the primary site is unavailable, the edge server will obtain the requested content from the mirror site in a fashion invisible to the end user.

**Failover Option 3—Failover to Edge Server:** If the edge server needs to contact the origin server to fetch or revalidate content, but cannot reach the origin server, it can be configured to serve the expired (most recent) version currently in cache. Agencies can configure the time it takes for the Qwest server to time-out its attempt to connect to the origin server and

serve the most recent content instead.  depicts the failover to edge server process.



The needs of a particular Agency's site and available infrastructure will determine which site failover option is appropriate. In all three scenarios, however, Qwest automatically detects whether the customer's origin server is responding to requests, and will detect when it is back online.

**Potential Service Enhancements:** Additional content delivery features, application accelerators, on-demand events, and performance management tools can be provided. These include:

- **Flexible Time to Live (TTL) and Time to Restore Settings:**  
Agencies can define the TTL for every object or page, a designation that can be assigned in less than a minute. TTL can be set from "no store" (never cache) to seconds, minutes, hours, or days. The Agency can also direct not to cache certain objects, or to check with the origin at every request for a particular object

to see if it has been modified. A page often consists of five to twenty objects, and each can have their own custom TTL.

- **Access Control:** Access control allows Agencies to limit access to content by integrating with authorization policies defined on an origin server.
- **Advanced Cache Control:** Advanced cache control enables [REDACTED] to increase the ability to store complex and dynamic content.
- **Content Targeting:** [REDACTED] content targeting enables Agencies to customize content to drive targeted business strategies online. The possible applications are limited only by the imagination.
- **Download Manager:** Download manager provides a simplified method of distributing, downloading, and installing digitized assets via the Internet. It can be used with websites that deliver content via SSL, as well as with sites that require authentication before providing access to content. Download manager is available as an add-on component for Agencies who use their websites to deliver digitized files such as software, movies, or other large objects.
- **Dynamic Content Assembly:** Dynamic content assembly enables Agencies to assemble and customize Web pages on [REDACTED] edge servers, delivering personalized content reliably to every user, while minimizing the demands on centralized application servers.

- **Enhanced DNS:** Qwest's Enhanced DNS provides a robust, reliable, and scalable solution to direct end users to customer Web sites. It requires no change to existing DNS administration processes, and provides unparalleled reliability, scalability, and performance of DNS resolution.
- **FTP:** Qwest's File Transfer Protocol (FTP) is a managed service that incorporates proprietary replication technology and global traffic management service using best-of-breed core storage equipment. The result is a scalable, high-performance, and highly available storage and FTP download service.

[REDACTED]

- **Secure Content Delivery:** SSL processing is extremely slow and often requires content providers to substantially over-provision sites to maintain performance and scalability [REDACTED] [REDACTED] is a highly secure solution that addresses the performance and security needs of Agency SSL content, while reducing costs and complexity. It supports the reliable and secure delivery of SSL objects and pages, and runs on a dedicated section of the Platform. [REDACTED] offers the highest degree of physical security and is optimized for SSL traffic.

[REDACTED]

- **Tiered Distribution:** Tiered distribution is offered specifically to enterprises that experience flash crowds, or that offer a large number of sizeable files for download. [REDACTED] enables Agencies to effectively and quickly deliver content to end users while minimizing the number of hits back to the origin Web site. With [REDACTED], Agencies ensure high performance and dependability for their end users while reducing their Information Technology (IT) staff's planning requirements and costs.

**4.2.7.6 Qwest Experience with CDNS Delivery (L.34.1.4.6(f))**

Qwest and [REDACTED] have collaborated to deliver CDNS to both commercial and Government customers [REDACTED]

[REDACTED]

[REDACTED] Together, we offer a coordinated, proven capability to the Networkx program.

[REDACTED] is deployed in 11 of 15 Government Cabinet Agencies, has more than 40 Federal contracts, and was recognized by a trade publication in 2005 for delivering the eight fastest Federal websites, demonstrating both market share and platform performance.

Qwest and [REDACTED] have extensive experience with and understanding of the Government environment. **Figure 4.2.7-4** summarizes our CDNS experience.

**Figure 4.2.7-4. Qwest Team CDNS Experience.**

[REDACTED]	[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]	[REDACTED]
[REDACTED]	[REDACTED]	[REDACTED]

Qwest’s award-winning Customer Care team is available 24x7x365 to answer Agencies’ questions and effectively respond to service issues.

Qwest will augment our Contractor Program Office (CPO) with the [REDACTED] Public Sector group to deliver CDNS. [REDACTED] large and experienced Public Sector team focuses exclusively on the Government market, and has a proven service delivery approach. [REDACTED] technical consulting team has supported the design of every CDNS deployment, and its professional services team has been responsible for more than 50 deployments in the last 36 months. Proven processes are in place,

experienced personnel are dedicated to the Government market, and a corporate support team will provide any additional support required.

**4.2.7.7 Characteristics and Performance of Access Arrangements  
(L.34.1.4.6(g))**

For the majority of Networkx services, Service Delivery Point (SDP)-to-SDP service quality depends on the facilities of more than one provider. Qwest realizes that a key factor to our success on the Networkx program is the ability to manage access arrangements from Agency locations to our core network through both the traditional Incumbent Local Exchange Carriers (ILECs) and the growing diversity of Competitive LECs (CLECs). Qwest has a unique combination of experience, as both an ILEC and as the first service provider to have successfully merged the operations of an ILEC with a long-distance company. This assures Agencies of our ability to provide robust access solutions that will meet Agency needs throughout the term of the Networkx program. Qwest uses the same discipline and approach that are used to maintain our own facilities-based portions of the service to the end-to-end delivery of Networkx services.

Details of CDNS access arrangements may be found in Section 3.2, Approach to Ensure Service Quality and Reliability, to include wireline and broadband access arrangements.

**4.2.7.8 Approach for Monitoring and Measuring CDNS KPIs and AQLs  
(L.34.1.4.6(h))**

The Qwest Control Networkx Portal will provide Agencies with easy-to-use monitoring and measurement tools to support management, operations, and billing of CDNS, including performance data. Through our portal, Agencies can view their traffic, content, applications, and users. Because our platform is open, Agencies can integrate the activation, management, and monitoring of services, content, and applications.

Qwest uses the following approach to measurement of CDNS Availability. This approach will also measure performance enhancement relative to the existing origin site, including Time to Refresh (GoS) and Time to Restore:

- From at least six geographically and network-diverse locations in major metropolitan areas, we will simultaneously poll a test file residing on the Agency's production servers and on [REDACTED]'s network.
- The polling mechanism performs two simultaneous http GET operations:
  - a. A test file is placed on the customer's origin server (i.e., origin.customer.com)
  - b. One GET operation is performed to retrieve the file directly from the origin server (i.e., "http://origin.customer.com/testobject")
  - c. The other GET operation is performed to retrieve the file through the service by requesting the object from the appropriate Agency hostname CNAMEd to [REDACTED] (i.e., "http://www.customer.com/ testobject" where "www.customer.com" is CNAMEd to [REDACTED] and configured to pull content from "origin.customer.com")
- The test content must use a Time To Live (TTL) of two hours or greater.
- The test content will be a file of approximately 10 Kbps in size.
- Polling will occur at approximately 12-minute intervals.
- [REDACTED] may also leverage third-party performance measurement providers such as Keynote Systems and Gomez Networks. These providers have performance measuring agents (i.e., servers that

simulate end-user activity) throughout the world. An example of a measurement may consist of downloading a test object directly from the origin site and the same object from [REDACTED]. Download times are recorded and presented.

- Based on the http GET operations, the response times received from the two sources—1. the Customer server directly; and 2. the [REDACTED] network—will be compared to measure performance metrics (latency, time to refresh), outages, and time to restore.

Each performance metric is based on a daily average of performance for the service, hits, and the Agency's production Web server—measured directly and computed from data captured across all regions. An outage is defined as a period of at least two consecutive failed attempts, six minutes apart, by a single agent to GET the Agency's test file from the service, while succeeding to GET the test file from the Agency origin server directly. In order to activate the conformance to the AQL, the Agency must enter and indicate the location of two valid test files for the same object into the portal exposed Service Level Agreement (SLA) Activation Tool. Detailed instructions are provided with the SLA Activation Tool. In addition, assistance is available from the Qwest Account Manager. The AQL will go into effect within five business days after the customer enters valid test files into the SLA Activation Tool.

Qwest's CDNS will establish a configuration file for each website or application. Within the GoS (Time to Restore) configuration file, refresh rates or caching times may be defined down to very granular levels (i.e., every object on every page.) Once the content refreshing rules are established, the content may be updated continually by the content manager. Agencies may view the configuration file and monitor performance using the Qwest Control Network Portal.


For all the services that Qwest offers, we use the [REDACTED] trouble ticketing system. [REDACTED] is a trouble ticketing system that is an industry-leading off-the-shelf commercial application that we have customized to make more effective for our needs. From this system, we collect many useful metrics that we use internally to evaluate and improve our processes including TTR. The calculation for TTR uses the same business rules as the Government requires for its services.

### **Measuring SDP-to-SDP Latency, and the Role of SEDs**

All of Qwest's IP-based services—which include Internet Protocol Services (IPS), Network-Based Internet Protocol Virtual Private Network Services (NBIP-VPNS), Premises-Based Internet Protocol Virtual Private Network Services (PBIP-VPNS), Layer 2 Virtual Private Network Service (L2VPNS), Converged IP Services (CIPS), CDNS, Voice Over Internet Protocol Transport Services (VOIPTS), and Internet Protocol Telephony Service (IPTelS)—are provided over the same IP services infrastructure. As a point of reference, Qwest has structured its network into a set of Provider (P) routers which form the core of our network and a set of Provider Edge (PE) routers that provide access to network services.

The P routers are private, in that they do not participate in routing to public Internet addresses. Their function is to provide bandwidth between the several sets of PE routers via sets of full-mesh Label Switch Paths (LSPs).

Following standard convention, the SDP is the Customer Edge (CE) router, as depicted in [REDACTED]



If an Agency orders a service in which the technical performance requirements are specified on an SDP-to-SDP basis (including performance requirements specified on an end-to-end and/or Agency premises-to-Agency premises performance requirement basis) and where Qwest requires the use of SEDs to meet the requirements and/or requires access to, or use of, the Agency's CPE or software to meet the requirements, then Qwest understands that the ordering Agency may (1) elect to not order such SEDs and/or (2) elect to not permit Qwest access to, or any use of, the Agency's customer-premises equipment or software for such purposes.

Qwest further understands that in this situation(s) and unless otherwise agreed to by Qwest and the user Agency, Qwest, when directed by the user Agency or by General Services Administration (GSA), will monitor, measure, and report the performance of the service for KPI/AQL and for SLA purposes either (1) on an SDP-to-SDP basis, by defining the SDP for performance metric measurement purposes for affected location(s) as being

located at the connecting POP(s) of the location(s), or (2) on a POP-to-POP basis. If directed to use the latter method by the Agency, Qwest will comply with the following:

1. For all IP-based network services, the applicable POP-to-POP performance requirements to be used will be those defined in Section C.2.4.1 (IPS).
2. For all other services, the service-specific SDP-to-SDP performance metrics will be applied on a POP-to-POP basis unless a stipulated POP-to-POP performance metric already applies for the associated service(s).

Qwest will provide three monitoring options:

- Standard SDP-to-SDP approach
- POP-to-POP approach
- Auxiliary SED for SDP-to-SDP monitoring approach

The standard monitoring for SLA reporting operates as follows:

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

Qwest's approach consumes few resources at the SDP (generally the CE router) as probes are sent from the Qwest network. This methodology does require that the customer respond to ICMP ECHO (a.k.a. ping) messages.

[REDACTED]

[REDACTED]

**The Use of Government Furnished Property (GFP)**

If an Agency orders a Transport/IP/optical service in which they are employing a GFP device, Qwest will provide KPI monitoring and measurement of the delivered service in three ways:

1. Request the Agency provide SNMP capability to the device for the Qwest NOC
2. Request that the Agency buy a monitoring SED from Qwest

3. Coordinate with the Agency for the following:

- Qwest understands that the ordering Agency may (1) elect to not order such SEDs and/or (2) elect to not permit Qwest access to, or any use of, the Agency's SED or software for such purposes.
- Qwest further understands that in these situation(s), and unless otherwise agreed to by Qwest and the Agency, Qwest, when directed by the Agency or by GSA, will monitor, measure, and report the performance of the service for KPI/AQL and for SLA purposes either (1) on an SDP-to-SDP basis, by defining the SDP for performance metric measurement purposes for affected location(s) as being located at the connecting POP(s) of the location(s), or (2) on a POP-to-POP basis.

**4.2.7.9 CDNS Support of Time-Sensitive Traffic (L.34.1.4.6 (i))**

All of Qwest's data networking solutions provide proven, industry-standard methods to ensure the quality of time-sensitive traffic. Qwest has best-in-class technical solutions and implementations of Quality of Service (QoS) mechanisms for both our integrated Asynchronous Transfer Mode (ATM)/Frame Relay (FR) network and our underlying Internet Protocol Multi-Protocol Label Switching (IP/MPLS) network.

Our homogeneous integrated ATM/FR network (which could provide access to CDNS) enables our customers to pre-assign applications to service classes for virtual circuits.

For ATM these are:

- Constant Bit Rate (CBR)
- Variable Bit Rate real-time (VBR-rt)
- Variable Bit Rate - non real-time (VBR-nrt)
- Unspecified Bit Rate (UBR)


For Frame Relay these are:

- Variable Frame Rate – real-time (VFR-rt)
- Variable Frame Rate – non real-time (VFR-nrt)
- Unspecified Frame Rate (UFR)

In general, voice is assigned ATM, CBR, or VBR-rt. Video is assigned VBR-rt, and other data traffic could be assigned UBR. The Qwest network completely conforms to ATM and FR standards, with the QoS enabled on a per-virtual circuit basis end-to-end, and strictly adheres to ATM and FR traffic contracts.

Traditional IP networks have evolved around “best effort” service, and typically have not provided guarantees for key performance criteria. The need to support real-time services on IP networks has driven the development of IP prioritization and queuing mechanisms, as well as MPLS technology. The Qwest network is engineered to enable QoS to prioritize certain types of traffic over other types of traffic if there is congestion in the network.

Qwest’s MPLS network supports the ability to prioritize Label Switch Paths (LSPs). This means that the IP network supporting our VoIP network has a higher priority than our VPN network, and that the IP network supporting our Voice over Internet Protocol (VoIP) network has a higher priority than the network that provides Internet services. Because of our ability to manage and prioritize traffic, impacts from different traffic loads can be handled immediately to ensure no impact to the bandwidth required to support all of our customers’ Virtual Private Network (VPN) and VoIP traffic requirements. [REDACTED] highlights the quality of service enabled by Qwest’s converged IP MPLS. Qwest’s IP MPLS network employs standards-based MPLS and IP-based QoS mechanisms to enable high quality voice,



video, and data over an IP backbone. The process of applying QoS in a network consists of multiple actions, defined as follows:

- Classification: Classifies different applications based on their relative network performance needs. For example, is the point-of-service application more or less sensitive to latency, jitter and loss than the VoIP application?
- Marking: Marks and classifies packets belonging to the applications so they may be recognized. For example, setting the Internet Protocol Precedence or DiffServ Code Point (DSCP) bits.
- Policing: Packets determined to be out of profile (that is, not conforming to the QoS policy), are either dropped or re-marked into lower priority packets (for example, rate-limiting).

- Shaping: Out-of-profile packets may also be buffered and shaped to conform to the configured QoS policy.
- Queuing: Scheduler resources are allocated to different classes (or queues) so traffic may be serviced (for example, last in, first out; first in, first out; weighted fair queuing; and low latency queuing).

[Redacted content]



[Redacted text block consisting of multiple lines of blacked-out content]


[REDACTED]

These QoS actions ensure that low latency, real-time applications such as voice can share the same access lines and core with non real-time data applications. Our convergence approach means that Qwest data services will migrate to a common IP/MPLS network, so we can easily plan and identify any quality of service issues. As described in Section 3.2.3, Qwest's backbone and access bandwidth planning methodology ensures that there is sufficient bandwidth to meet an Agency's full port-limited capability, even in the event of core router failure, or an access router or backbone trunk failure.

**4.2.7.10 CDNS Support for Integrated Access (L.34.1.4.6(j))**

Qwest's network architecture and data services approach directly enables a complete menu of integrated access options to virtually all of Qwest's services.

As depicted in [REDACTED] Qwest's network architecture and services approach offers a broad range of already integrated access methods:



[Redacted]

[Redacted]

[Redacted]

[Redacted]

[Redacted]

[Redacted]

[Redacted]

[Redacted]

[Redacted text block]

**4.2.7.11 Infrastructure Enhancements and Emerging Services  
(L.34.1.4.6(k))**

Qwest has mature processes that enable us to envision, research, evaluate, engineer, deploy, and operate new or emerging services. Driven initially by the Chief Technology Office, headed by the Qwest Chief Technology Officer (CTO), Qwest evaluates new products and technologies for incorporation into the Qwest network, in partnership with Qwest Product Management.

[Redacted text block]

[Redacted text block]

[Redacted text block]

[Redacted text block]

Qwest recognizes that converged customer care and support will be a major challenge that impacts processes, systems and people. Convergence extends and impacts every facet of the traditional telecommunications value chain.


[Redacted text block]




**4.2.7.12 Approach for Network Convergence (L.34.1.4.1(I))**

Qwest already has a clear approach and has made significant progress in deploying a network that not only enables convergence from the customer's perspective, but is also a highly converged platform in itself. Qwest is moving toward a packet-based architecture to enable network evolution and convergence. Centered on our private MPLS-based core, Qwest has already converged our IP-based services (private port iQ MPLS VPNs, public port iQ for Internet services, and our VoIP transport for Public Switched Telephone Network (PSTN) traffic) over this network.

Qwest is committed to the elimination of stovepiped networks that create planning, operations, and interoperability issues for our customers.


 shows Qwest's approach to ensure that services have a uniform view of network and support infrastructure.




Multiple overlay networks are no longer established to deliver new services. Value is shifted to network-based services, where Qwest becomes a solutions provider. Applications-based services are delivered independent of the network infrastructure. Excellent service quality is maintained during network convergence through the following practices:

- Consistent and rigorous technology management methodology that includes evaluation, selection and certification of network elements
- Accommodation of legacy services as the network evolves
- Network simplification through de-layering and introduction of multi-service access devices
- Coincident convergence of back-office systems, including introduction of a next-generation network management layer packet operational support system

As shown in [REDACTED], the use of a converged MPLS core significantly eases the problems normally associated with backbone traffic engineering. Without a converged backbone, each services network (for example, one for Internet, one for private IP services, and one for Voice) needs to be traffic engineered independently. The normal state of affairs is that one network has too much capacity and another has performance limitations that require a backbone or router upgrade. The issue is that a carrier gets into a situation where the upgrade for one services network requires a large upgrade that is not cost effective. For example, the desired upgrade from OC-48c to OC-192c backbone circuits may require a complete nationwide upgrade that the carrier cannot afford, forcing them to settle for suboptimal performance regarding SLA fulfillment.



**4.2.7.13 IP-PSTN Interoperability (L.34.1.4.6(m))**



[Redacted content]

**4.2.7.14 Approach for IPv4 to IPv6 Migration (L.34.1.4.6(n))**

Qwest is well positioned to migrate its network from IPv4 to IPv6.

[Redacted text block consisting of 10 horizontal bars]

[Large redacted text block]

[Redacted text block]

[Redacted text block]

**4.2.7.15 Satisfaction of NS/EP Requirements (L.34.1.4.6(o))**

According to RFP Section C.5.2.2.1, NS/EP Basic Functional Requirements Matrix for Network Services, CDNS is not NS/EP required. Details of how Qwest supports the 14 basic functional requirements for applicable services are provided in Section 3.5.1, Approach to Satisfy NS/EP Functional Requirements, in this Technical Volume.

**4.2.7.16 Support for Signaling and Command Links (L.34.1.4.6(p))**

[Redacted text block]

[REDACTED]

**4.2.7.17 Service Assurance in the National Capital Region (L.34.1.4.6(q))**

As discussed in Section 3.2, *Approach to Ensure Service Quality and Reliability*, Qwest provides network services in the National Capital Region (NCR) with a robust network architecture designed and engineered to ensure service continuity in the event of significant facility failures or catastrophic impact. Qwest will continue to engineer critical services to meet each Agency's requirements to eliminate potential single points of failure or overload conditions that may impact their network service performance.

[REDACTED]



Qwest has fully described our approach to satisfying Section 508 requirements for applicable, offered services in Section 3.5.4, *Approach for Meeting Section 508 Provisions*, of this Technical Volume.

**4.2.7.19 CDNS Impact on Network Architecture (L.34.1.4.6(s))**

Qwest anticipates no impact to our network architecture due to the delivery of CDNS.

**4.2.7.20 Optimizing the Engineering of CDNS (L.34.1.4.6(t))**

Qwest's CDNS platform optimizes content assembly and delivery to efficiently deliver Web content and applications from carefully located and load-balanced servers. The content delivery platform is connected with the clients' content generation infrastructure using optimal paths through the Internet, and using intelligent routing algorithms supported by real-time network information. The delivery of content is then served by [REDACTED] edge servers that are deployed near all end users.

**Intelligent Request Routing:** Qwest's CDNS platform uses patented Internet monitoring software to maintain a comprehensive view of network health, and sophisticated content-routing algorithms to route users to the optimal servers on the platform.

We employ a variety of techniques for Internet topology discovery, and for measurement of up-to-the-minute latency, loss, and bandwidth metrics to a variety of points on the Internet. These network performance measurements are then combined with detailed load, liveness, and capacity information from our servers to arrive at mapping decisions for end users. As a result, every end user is mapped to a nearby server that is lightly loaded, and that maximizes performance for the relevant application. Manifested as an instantaneous decision at DNS resolution time, our patented network-routing algorithms find the best edge server for each request. This ability to optimize load and performance is unique to our CDNS platform.

**High Performance Communications:** The Qwest CDNS platform has highly optimized communications between edge servers, as well as between edge and origin servers, to ensure that content and data are always readily available from all edge servers. [REDACTED]

[REDACTED] identifies alternate paths from our edge server to the origin server, and uses those alternatives to improve performance of content delivery. [REDACTED]

[REDACTED] is a highly scalable and reliable system used to transmit Agency configurations (metadata) to edge servers. Unique aspects of this infrastructure are its application of [REDACTED] for high-performance routing to edge servers, robust mechanisms for replicating data submitted to our content delivery platform, and the ability to transport data to the entire content delivery network very rapidly.

**Network Management and Monitoring:** To ensure ongoing optimization of this diverse and distributed platform, Qwest has built a comprehensive set of tools to administer a network configuration and heterogeneous state efficiently, to enable easy modular modification and addition of software components, and to scale and fully automate the software installation process. The Network Operations Center (NOC) uses proprietary, secure, scalable, real-time data collection mechanisms to enable efficient and responsive monitoring of the platform. If a problem is detected, our fault-tolerant architecture takes over, automatically switching from one edge server to another. Our NOC personnel investigate the cause.

**Proactive Performance Monitoring:** In addition to the NOC, Qwest has a variety of software and infrastructure that provides detailed information used to ensure optimal performance for end users, as well as for continual analysis and optimization of the CDNS algorithms and network. Examples include data collected from [REDACTED] distributed agents for HTTP and HTTPS testing, and [REDACTED] proprietary, patented agents and technology for testing streams from all major formats, and network protocol-level statistics logged by [REDACTED] servers.

**4.2.7.21 Vision for Service Internetworking (L.34.1.4.6(u))**

Qwest, the leader in the development of IP-based convergence, and [REDACTED], the world's leading IP delivery vehicle, firmly believe in a common, IP-centric architecture. Qwest and [REDACTED] have implemented to this vision. The foundation of the content delivery business lies in providing IP-based services that leverage common standards and network interoperability. Qwest's broad and deep solution set of IP-based service offerings ensures support for any CDNS or related requirement, ranging from the division to the department. Qwest is well positioned to support the continuing evolution toward IP convergence, delivering content around the world, to meet the ever-increasing demand for IP-based services.

**4.2.7.22 Support for Government CDNS Traffic (L.34.1.4.6 (v))**

Over the course of the 10 years delineated in the Government's pricing model, committed bandwidth for both domestic and global, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]