A Practical Review of Broadband Requirements For Healthcare Clinical Applications

November 2009

Prepared by:

Fiberutilities Group, LLC
222 Third Street, SE, Suite 500
Cedar Rapids, IA 52401
(319) 364-3200
www.fiberutilities.com

INTRODUCTION

Fiberutilities Group, LLC is a professional utilities management company, providing communications connectivity management, planning, building and operating services to clients. Fiberutilities Group currently provides full operational control of over 8,000 miles of fiber optic network including first/last mile, middle mile and long haul backbone connectivity at the operational, data link level, including monitoring and maintenance of the physical network, as well as providing design, planning and construction services for any new fiber network required. For those clients with higher levels of operational sophistication Fiberutilities Group provides incremental support services at every layer starting with acquiring network assets (dark fiber) and moving up the value/service chain to full network operation and control.

Fiberutilities Group has worked with several healthcare systems across the central USA in researching available options and procuring last mile and middle mile network connections (i.e. WAN circuits) to support the operation of specialty, urban, and rural clinics as well as creating private healthcare networks to interconnect with larger networks in large metro markets.

For this document, Fiberutilities Group assumes that “broadband” is a descriptive term for one type of clinic access connection that can promote more widespread adoption of healthcare applications important to the efficacy of federal healthcare goals. The focus of the discussion will then be to examine the capacity, quality, and characteristics of broadband connections that can enable these higher goals.

ARCHITECTURE & NETWORK TOPOLOGY

Health systems that have developed communications networks for their own use have one characteristic that allows them to operate as collaborative units. Each group has specific business needs it is addressing, and each business need has a specific application that provides the technical capabilities. In such an environment, network development is relatively simple: connections from one point to another have a specific purpose and a specific A and Z location. These health systems are primarily focused on operational cost savings delivered by shared resource consumption including such things as common IT applications and common medical technology solutions. This “packaged” environment is easier to deal with than an environment in which no data centers, application services, or specific service platforms exist. Said differently, the problems of Electronic Health Records (EHR) and other medical solutions may be harder to solve as a general problem because the communications focus is driven toward the access end of the system. The primary issue with EHR and other telemedicine proposals is that the applications themselves have no predefined network architecture until a specific vendor solution is chosen. Once a vendor is chosen then a network can be designed to address the specific communications needs of the medical group.

To deal with the structure of healthcare organizations and desired connectivity, a logical starting point is the creation of data application “Malls” or storage centers/depots in the core of the medical network. This infrastructure would then become the data connectivity destination for medical connectivity. Having provided a functional topology for medical applications, the size requirement for the access end can then be resolved on an application basis. Broad discussions about latency and data throughput can be greatly simplified by assuming an application architecture that regionalizes the basic medical applications. This regional topology, with an interconnected middle mile will create a data and application matrix on which vendors can provision their medical applications using similar access methods connected to service and data depots. Depots can be interconnected to provide extra-regional access. This approach could enable the further development of Internet2, National Lambda Rail or a similar, healthcare-specific network by focusing financial resources on a core problem, that of medical application accessibility. It would also enable the medical community to move away from the Internet as a service access platform.

LATENCY AND DATA RATE

Once viable application architecture has been developed, the problems of latency (request-response delay), data rate (volume), and jitter (packet delay inconsistency) can be addressed as required by each application. A technical transaction in communications systems is often an accumulation of requests and responses that result in the assembly of a presentation to the user. Such a presentation can be optimized by reducing either the number of bits transferred, the number of transactions required, or by raising the priority of the application in the communications pool. Three common approaches are terminal emulation, packet expansion, and QoS prioritization. Terminal emulation uses a decreased amount of bit transfer to optimize applications by attempting to process data in a central server to the greatest extent possible prior to sending responses to the client. For a terminal emulation approach increased latency will cause application delays because the application has a high ratio of transactions to data volume. For a PACS transaction, larger packet sizes are used, and larger data requests are made thus reducing the ratio of transactions to data. This sort of application is susceptible (i.e., can degrade) because of low data rates, but is less susceptible to high latency. A video transaction requires constant feedback thus creating a high data rate and a low latency need. Increasing the priority of this type of data packet over other packets in the same stream best optimizes this sort of application.

Once viable application architecture is proposed, then applications can be deployed on the architecture using any of the above approaches to best suit the application. The problem will become less an access issue, and more controllable as a function of the architecture used to make medical applications available.

CLINICAL ACCESS NETWORK

The threshold characteristic for clinical broadband connectivity is that the clinic is likely connecting to many destinations, and each application may have a different destination and different set of network requirements. For example, if a clinic uses an email server other than a centralized corporate server, clinic computers must connect to a server somewhere (e.g., anywhere on the Internet) to send and retrieve messages. The time required for the request data packet and the response data packet is not particularly important. The size of the connection may be a more important factor in the speed with which asynchronous packets can get through the connection.

If, for example, the clinic is using a terminal emulation-based enterprise clinical system to manage financial and administrative functions, the connection is to a specific host application likely behind a secure firewall. Moreover, the application may be looking up records in a database, thus adding a response requirement (measured in round-trip packet travel time, or latency) on the connection to support real-time applications. Latency can be more important than bandwidth for these applications to perform properly. So while bandwidth is of key importance for certain types of applications (email, web surfing), latency may be the most critical factor in providing adequate connections for centralized database applications (EHR/EMR and clinical systems).[1] Traditional Time-Division Multiplexed (TDM) circuits like DS-1's offer low latency because they have few hops involved in the transmission. The Internet, on the other hand, is comprised of many router hops as a means for networks to interconnect to route packets to the proper destination. This difference is illustrated in the Figure 1 below. The trade-off: while traditional TDM circuits offer low latency, they also offer less bandwidth. The Internet (through broadband access technologies like DSL and cable modems) can offer greater bandwidth speeds but the connected networks architecture can significantly increase latency – especially in more remote areas that require several layers of ISP providers to route the traffic to the destination.

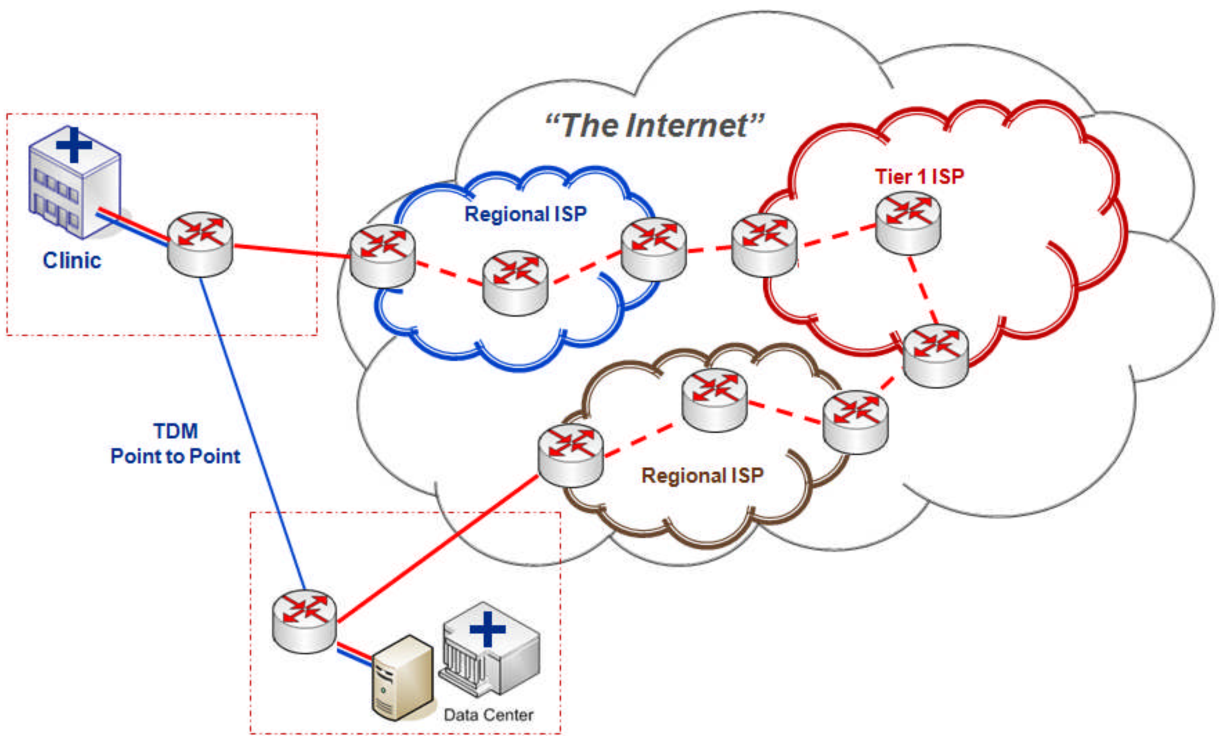

Figure 1

To document this point, below is a trace route of data packets from a rural location in Iowa to an Iowa healthcare provider public website. The data packets passed through 11 router hops, each hop adding latency to the path.

```
Traceroute to xxx.org from rural Iowa
11 packets transmitted, 11 packets received, 0% packet loss
round-trip min/avg/max = 65.031/64.777/64.765
```

[1] If the network link has low bandwidth then it's an easy matter of putting several in parallel to make a combined link with higher bandwidth, but for a network link with bad latency then no amount of money can turn any number of them into a link with good latency. http://rescomp.stanford.edu/~cheshire/rants/Latency.html

```
traceroute to xxx.org (205.167.2.50), 64 hops max, 40 byte packets
 1  192.168.0.1                          (192.168.0.1)      0.763 ms    0.456 ms    0.273 ms
 2  cdrr-dsl-gw01-193.cdrr.qwest.net     (209.181.206.193)  39.246 ms   39.327 ms   40.227 ms
 3  209-181-211-1.cdrr.qwest.net         (209.181.211.1)    41.891 ms   41.339 ms   40.338 ms
 4  cdr2-core-01.inet.qwest.net          (205.171.163.33)   41.150 ms   41.049 ms   39.628 ms
 5  cdr2-core-02.inet.qwest.net          (205.171.163.2)    39.811 ms   39.600 ms   40.607 ms
 6  cer-core-02.inet.qwest.net           (205.171.5.94)     45.437 ms   43.945 ms   45.539 ms
 7  chp-brdr-03.inet.qwest.net           (67.14.8.190)      45.815 ms   44.351 ms   45.054 ms
 8  63.146.27.18                         (63.146.27.18)     45.761 ms
    63.146.27.22                         (63.146.27.22)     44.685 ms
    63.146.27.18                         (63.146.27.18)     45.273 ms
 9  ae-21-54.car1.Chicago1.Level3.net    (4.68.101.98)      45.901 ms
    ae-21-56.car1.Chicago1.Level3.net    (4.68.101.162)     45.300 ms
    ae-11-55.car1.Chicago1.Level3.net    (4.68.101.130)     45.678 ms
10  XXX.car1.Chicago1.Level3.net         (4.79.209.178)     50.617 ms   48.413 ms   49.992 ms
11  205.167.3.198                        (205.167.3.198)    65.031 ms   64.777 ms   64.765 ms
```

The roundtrip latency for this path is between 64 and 65 milliseconds. By contrast, a typical frame relay clinic circuit and the typical clinic T1 from a rural location to a central location are point-to-point circuits, and have typical response times of 30ms (frame relay) and 10ms (T1). This trace route is typical of the multiple hops necessary to deliver data packets outside of major metro areas. This particular example, however, is probably conservative because this health system has its own direct fiber connection to Chicago. That is not the case for most non-metro healthcare providers.

What all this implies is the need for new middle mile infrastructure that can combine the low latency characteristics of legacy circuit-based networks while also offering the high bandwidth, low and low cost characteristics of modern packet-based networks. This is one of the primary drivers for the creation of private middle mile network solutions in the healthcare industry.

NETWORK TESTING RESULTS

The following table summarizes the results of Fiberutilities Group's testing of available connectivity alternatives for clinical applications:

Summary Results:

Connection Technology	Bandwidth Expected	Latency	Response Time	Connection Type
Broadband DSL	768 Kbps / 256 Kbps	86 ms	1.3 Seconds	Point to Internet
Broadband MPLS over ADSL	768 Kbps / 256 Kbps	57 ms	1.9 Seconds	Point to Point
Business Cable Modem	12 Mbps / 2 Mbps	68 ms	1.8 Seconds	Point to Internet
Business Cable Modem *	12 Mbps / 2 Mbps	29 ms	1.4 Seconds	Point to Internet
Licensed P-to-P Wireless	100 Mbps / 100 Mbps	18 ms	1.0 Seconds	Point to Point
Fiber Metro Ethernet	100 Mbps / 100 Mbps	4 ms	.6 Seconds	Point to Point

* *This second test was run after working with the providers on middle mile engineering adjustments to keep traffic local (i.e. hair-pinning within the region). This is possible for certain end locations but probably not all locations because of IP scheme changes required.*

CONCLUSION

Based on Fiberutilities Group's experience and the expressly stated needs and concerns of its clients in the healthcare industry the "appropriate" amount of bandwidth for hospitals, clinics, and physicians is really dependent upon where the data is captured and where the data is needed as well as on the specific applications deployed. See Appendix A for a summary of typical health system network requirements. The current trends indicate rapid growth in the *amount* of data as well as in the need to *move* data where it is needed. This requires planning effort to consider the following factors:

- Assuming middle mile network exists and is available (it often isn't), adjustments can be made to improve the quality (and capacity) of the local loop access technology.

- While "bits are bits", application-specific requirements are impacted by more than raw capacity, both from a user-expectation perspective and an actual performance perspective.

- Availability drives bandwidth "demand," but the converse is not necessarily true. There is no demand for products that do not exist or exist at price points that exceed what the user deems reasonable.

- Current applications are tailored to the limited bandwidth available and to other network parameters such as latency and resiliency. An application designed for a TDM T1 connection may not work properly with a multi-Mbps or Gbps connection.

- The Internet is generally not sufficiently reliable, secure or of sufficient network "quality" for many healthcare related applications beyond simple email and online research.

- Private networks such as National Lambda Rail and Internet2 have the potential to provide better quality connectivity because they generally have fewer hops and greater backbone capacity dedicated to their users.

APPENDIX A

TYPICAL HEALTHCARE NETWORK REQUIREMENTS

These requirements assume the typical applications in use at a clinic may be any of the following:

- Email (sized based on the number of workstations)
- Internet access for research and education
- Electronic Health Records/Electronic Medical Records systems, depending on the software modules implemented
- Enterprise Clinical Systems used to manage clinical, financial and administrative functions
- Picture Archiving and Communication System (PACS) and other remote radiology applications capturing images

Based on these applications, the following are the minimum requirements for a rural healthcare access network:

- Standards-based (allowing interconnection with carriers) with QoS capabilities and standard network management tools
- Allowance for incremental growth as demand increases
- Point-to-point architecture with redundancy, and allow moves, adds and changes without network disruption
- Low latency (less than 50 ms primary, 120 ms backup/secondary)

SLA Targets

The following SLA targets are typical for the types of WAN connections used in clinical applications:

Connectivity Type	SLA Uptime (expected)	SLA Latency (expected)	SLA Bandwidth (available)	Restoration Time (in the event of failure)
Fiber WAN	99.999%	<20 ms	100Mb/s	<1 sec
Metro Ethernet	99.999%	<40 ms	10-100Mb/s	<1 sec
Point to Point T1	99.99%	<60 ms	1.5 – 6Mb/s	<1 sec
Broadband VPN	99.90%	<120 ms	1 Mb/sec	N/A

When considering how much bandwidth is required for a given location, the following general guidelines are appropriate:

- Workstations < 15: 1.5Mb/s or less; Workstations >15: 1.5Mb/s or greater

- Email – its effect is counted in the number of workstations
- Electronic Health Record System – the same formula for workstations can be applied to EHR systems: < 15 allows for 1.5Mb/s or less; Workstations >15: require 1.5Mb/s or greater
- Enterprise Clinical System – if it is simple terminal emulation it doesn't require significant bandwidth and its effect can be counted in the number of workstations
- PACS – Requires 100Mb/s to work well. Fiberutilities Group has run it on bonded T1's but eventually needed to upgrade to more bandwidth because it was operationally degraded (not fully functional)