DOD NETWORK TRAFFIC USAGE

Tony Modelfino, George Case
Stratogis Networks LLC
Potomac, MD

ABSTRACT

The design requirements for any network are driven by the circuits that will eventually connect to that network and the traffic they will support. In the past, the capacity requirements for DoD satellite networks have been based on point-to-point circuits. This was a reasonable approach given the dedicated bandwidth operational mode that was typical of these systems. For these networks, circuits would be provisioned and would tend to remain in place for long periods of time, regardless of the actual traffic on a given circuit. This service concept tends to be inefficient in its use of capacity, inflexible in its ability to adapt to changing traffic patterns, and limited in its ability to provide high-speed connectivity with guaranteed quality-of-service to all network nodes. New satellite payload and VSAT terminal technologies are integrating bandwidth-on-demand (BoD) and IP Quality-of-Service (QoS) into the network architecture, enabling very high speed any-to-any IP connectivity with service level guarantees. This paper describes a large-scale network monitoring effort that was undertaken to characterize DoD network traffic in an effort to assess the effectiveness of these new technologies. The authors analyzed the circuit usage characteristics of over 1500 NIPRNET and SIPRNET ports in the CENTCOM, EUCOM and PACOM AOR's during a 14-month period that included the Operation Iraqi Freedom ground campaign. This paper summarizes CENTCOM circuit utilization, provisioning accuracy, and circuit burtsiness on multiple time scales. Finally, we provide an estimate of the bandwidth utilization efficiency increase that can be achieved through the implementation of transformational technologies.

INTRODUCTION

Network data traffic is inherently bursty in nature. Network planners recognize the bursty nature of data traffic and have always over-provisioned circuits to meet peak or near-peak demands. The level of over-provisioning is often driven as much by economic factors as by technical considerations. In metropolitan areas where fiber is plentiful and bandwidth costs are low, network planners will tend to provision more bandwidth for a given demand. In isolated areas were bandwidth is scarce and costs are higher, provisioning will be more restrained. This paper will show that average circuit usage can increase two to three fold during combat operations. Thus, to evaluate the effectiveness of IP QoS Bandwidth-on-Demand (BoD) technologies for the DoD, it is critical to look beyond average circuit bandwidth usage and evaluate the both circuit and aggregate network burstiness under war time conditions. Since BoD networks employ dynamic resource sharing between nodes, aggregate bandwidth usage is the key metric for assessing acceptable operating loads for these networks.

The objective of this paper is to assess the advantage, or statistical multiplexing (STATMUX) gain, associated with a BoD architecture as compared to a traditional point-to-point dedicated bandwidth architecture. Other papers have used network simulations to estimate STATMUX gains by comparing the theoretical performance of BoD and static bandwidth networks as a function of network operating load. This paper estimates the STATMUX gain based on CENTCOM bandwidth usage during Operation Iraqi Freedom.

DOD NETWORK MONITORING

DISA (Defense Information Systems Agency) has instrumented and continuously monitors thousands of Tier-0 and Tier-1 NIPRNET and SIPRNET routers using NetHealth software to measure circuit usage, availability, latency and other parameters. The software records these statistics for each circuit every 5 minutes. The authors were given access to NIPRNET and SIPRNET circuit usage statistics from the PACOM, EUCOM and CENTCOM networks. Over 1500 network ports were monitored for a period of fourteen months including Operation Iraqi Freedom. These links included WAN connections such serial point-to-point and ATM circuits as well as Ethernet LAN connections. Our primary focus was WAN connections, which ranged from 64 kbps to 24 Mbps, since these were most representative of MILSATCOM circuits.

The authors received daily NetFlow log files for each network. These log files contained the average port usage over 5 minute intervals for each router in the network. The raw port data was processed as shown in Figure 1. During the course of the data collection effort, a number of circuits were added, deleted or upgraded particularly in the CENTCOM AOR. In an attempt to insure that usage characteristics were statistically significant, ports with less than 30 days of traffic were eliminated. An exception to this filter was a seven day minimum that was used for all OIF ground campaign processing, since the entire ground campaign period was only 30 days duration. This filter eliminated less than 10% of the circuits. A number of ports had very low usage or no usage at all. These may have been backup circuits or circuits that were no longer in use. These sites were clearly not representative of DoD traffic and if included in the analysis would substantially lower the overall port utilization; therefore, the authors eliminated any circuit whose average busy hour usage was below 1%. The minimum 1% usage requirement eliminated ~10% of the circuits. Average, 90^{th}, 95^{th} and 98^{th} percentile utilization characteristics were determined for each remaining circuit.

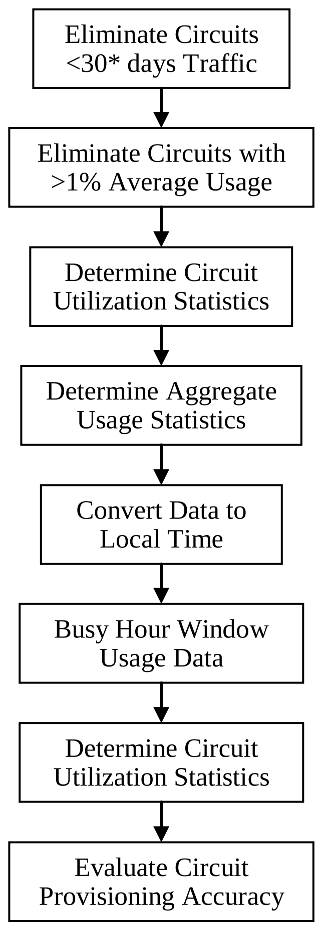

* 7 day filter used for processing during OIF

Figure 1 – Data Analysis Process

The data was then converted from ZULU time to local (Iraq) time and reprocessed to determine 8 AM to 6 PM circuit usage. The same metrics were determined for each circuit. As expected, the CENTCOM 8 AM to 6 PM performance characteristics do not differ significantly from the 24 x 7 characteristics during ground combat operations. The 8 AM to 6 PM usage in PACOM and EUCOM was significantly (200%) higher than the 24 x 7 usage, which is also expected since most nodes in these AORs were not principally involved in major combat operations in 2003.

The analysis sequence shown in Figure 1 was done for all PACCOM, EUCOM and CENTCOM circuits that were monitored. This paper will focus on CENTCOM AOR usage since we were interested in bandwidth usage during combat operations, when usage levels are the greatest. Our SIPRNET analysis focuses on the period from March 15 through April 15, 2003 when there was a surge in SIPRNET usage associated with combat operations. Figure 2 shows the daily traffic that was carried on the CENTCOM SIPRNET and NIPRNET circuits that were monitored. The data has been processed with a seven day averaging window. The peak in SIPRNET traffic occurred on March 31, 2003. As major combat operation concluded in late April 2003, SIPRNET traffic on the circuits monitored dropped by 45%. NIPRNET traffic, on the other hand, continued to increase as more personnel were deployed in theater. Analysis of NIPRNET traffic in early 2003 indicated that 75% of the traffic was Internet traffic as opposed to .mil NIPRNET traffic. Note that NIPRNET data was not available for the last three weeks in August 2003.

Tables 1 and 2 summarize the average CENTCOM SIPRNET and NIPRNET usage from March 15 through April 15, 2003 as a function of circuit speed. The average 8 AM – 6 PM usage was 28% for SIPRNET and 23% for NIPRNET. The lower circuit speeds tended to have the lowest utilization. This is expected given that low speed circuits tend to have burstier traffic characteristics and typically need greater over-provisioning to meet peak demands. There is a noticeable increase in utilization for fractional T1 circuits between 384 and 768 kbps on both the NIPRNET and SIPRNET. In both instances this increase can be attributed to a small number of circuits with average utilization in excess of 70%. Finally, there is little difference between the 8 AM- 6 PM and the 24 x 7 usage, owing to the round-the-clock nature of active theater operations.

Table 3 summarizes the NIPRNET usage over the full 14 month period. These results include significantly more circuits over a much longer period (400 days). The results show an almost constant average utilization of 24% across all circuit speeds. The results for the small number of circuits above 2 Mbps are slightly lower. Once again there is very little difference between the 24 x 7 and 8AM – 6PM results.

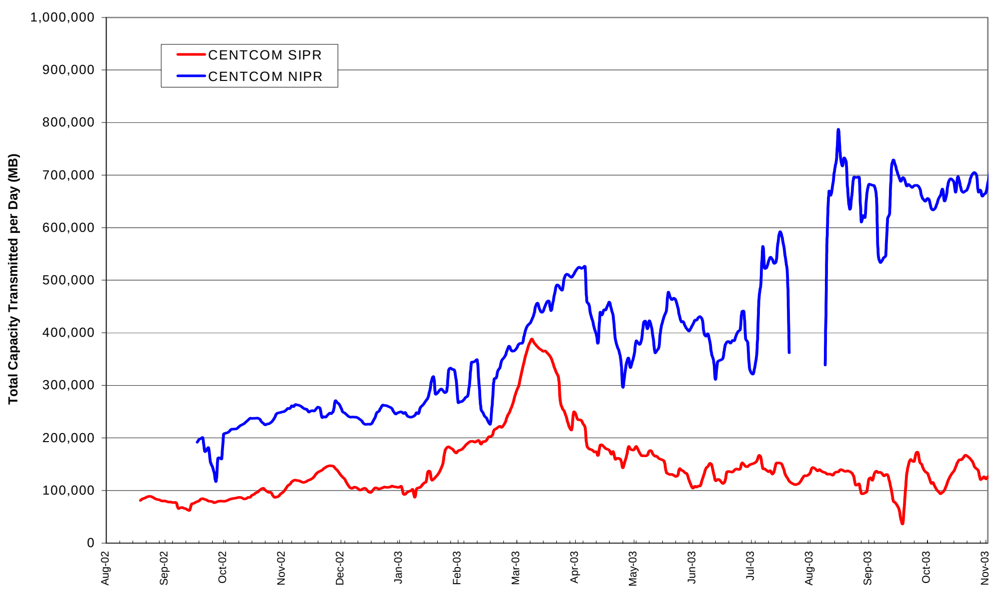

Figure 2 – Daily CENTCOM SIPRNET and NIPRNET Traffic

Dedicated point-to-point circuits can not be provisioned to meet the average demand, but instead must be provisioned to meet the peak or near peak bandwidth requirements. Figures 3 and 4 show the average, 90^{th}, 95^{th} and 98^{th} percentile NIPRNET and SIPRNET bandwidth utilization for CENTCOM circuit during OIF. The results show that between 2% (98^{th} percentile) to 5% (95^{th} percentile) of the time circuits are exceeding 50% utilization. Planners are clearly provisioning circuits to meet this peak demand. Since 5% corresponds to 1.2 hours per day or 1.5 days per month, planners must over provision to meet this demand.

Table 1 – Average CENTCOM SIPRNET Circuit OIF Utilization (3/15-4/15/2003)

Bandwidth (kbps)	Ports	8 AM - 6 PM	24 x 7
64-128	2	7%	5%
129-256	4	17%	18%
257-512	12	14%	13%
513-1024	7	37%	37%
1025-2048	42	25%	24%
2049-4096	18	20%	19%
>4096	2	34%	31%
Average All	87	23%	22%
Weighted Average	87	23%	22%

Table 2 – Average CENTCOM NIPRNET Circuit OIF Utilization (3/15-4/15/2003)

Bandwidth (kbps)	Ports	8 AM - 6 PM	24 x 7
64-128	10	10%	8%
129-256	14	19%	17%
257-512	14	47%	45%
513-1024	4	22%	20%
1025-2048	29	26%	24%
2049-4096	13	14%	13%
>4096	14	33%	28%
Average All	98	26%	23%
Weighted Average	98	28%	24%

Table 3 – Average CENTCOM NIPRNET Circuit Utilization (10/1/2002-12/1/2003)

Bandwidth (kbps)	Ports	8 AM - 6 PM	24 x 7
64-128	41	26%	26%
129-256	48	26%	26%
257-512	38	24%	24%
513-1024	32	25%	25%
1025-2048	66	23%	21%
>2048	13	16%	18%
Average All	238	24%	24%
Weighted Average	238	22%	22%

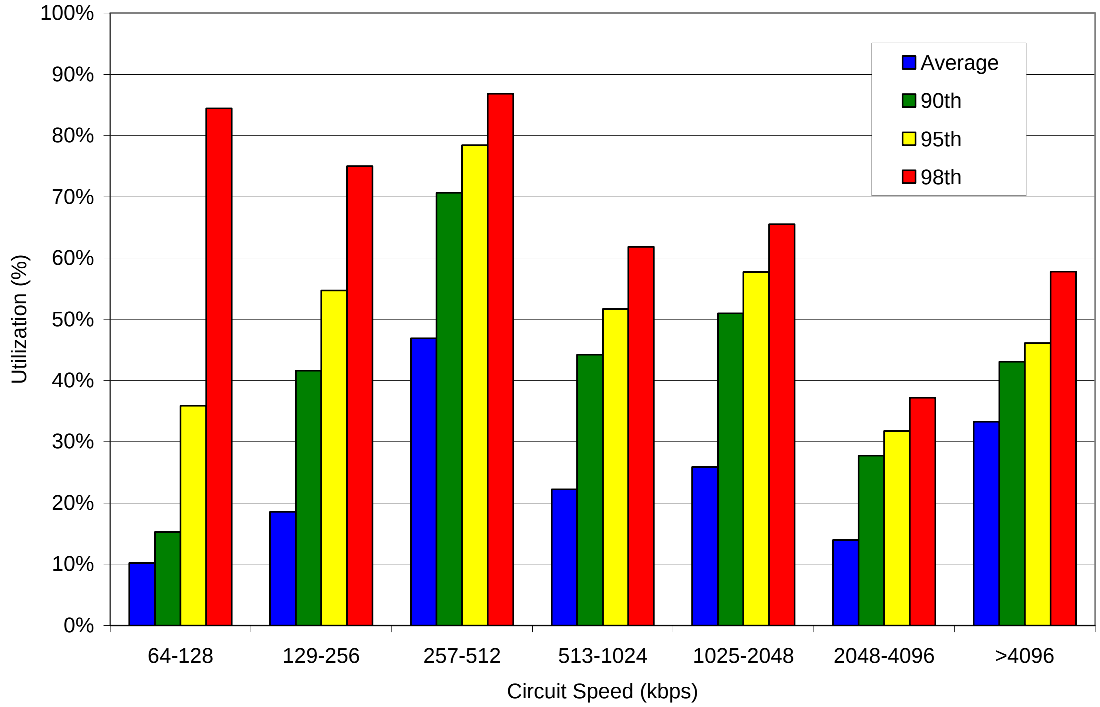

Figure 3 – NIPRNET Traffic Burstiness

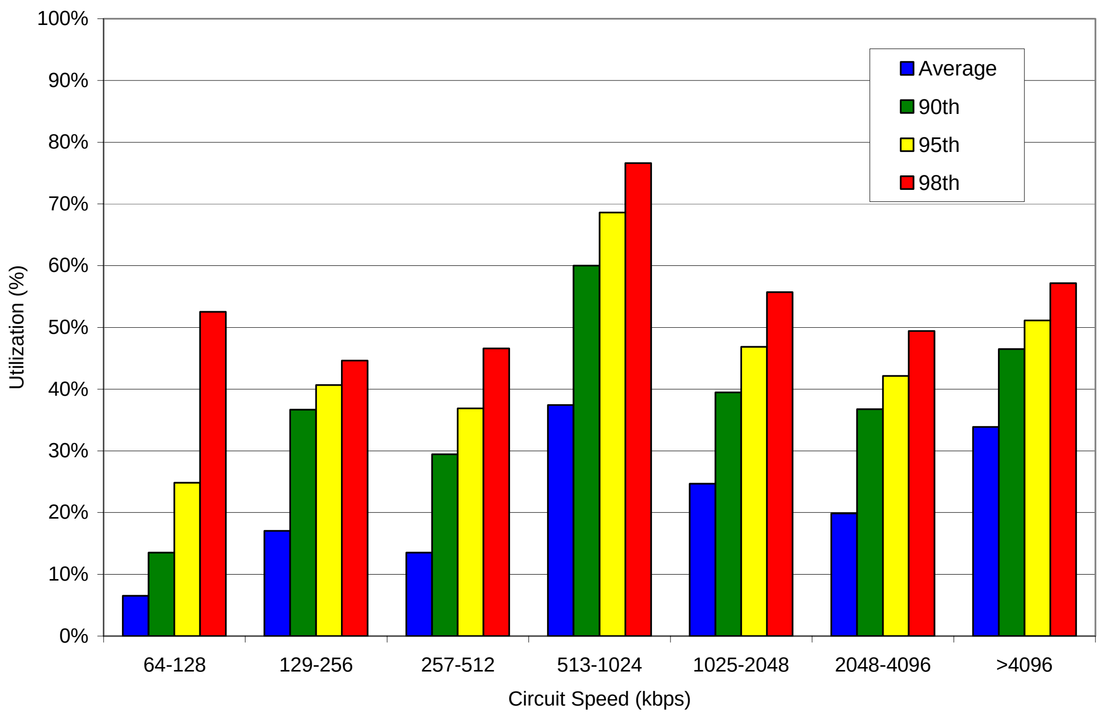

Figure 4 – SIPRNET Traffic Burstiness

Though the average utilization is 25%, a number of circuits on each network were highly congested. Table 4 shows the percentage of circuits that exceed 50% to 80% average utilization on each network. While there is no universal standard congestion threshold, the measurement indicates that between 2% and 16% of the circuits were congested and would have benefited from additional bandwidth.

Table 4 – Percentage of Circuits Exceeding Average Utilization Thresholds

Network	Average Utilization			
	50%	60%	70%	80%
SIPRNET OIF	10%	7%	2%	0%
NIPRNET OIF	16%	12%	7%	4%
NIPRNET	15%	13%	10%	5%

In an effort to evaluate the accuracy of circuit provisioning, we determined the percentage of circuits with 98th percentile utilization that exceeded 30% to 60%. The results are shown in Table 5. Since circuit bandwidth is typically only available at certain discrete rates; for example, 64, 128, 256, 384, 512, kbps etc, it would be impossible to provision the exact right bandwidth to meet the average and 98th percentile demand even if it was known in advance. If a network planner expected an average utilization of 64 kbps and a 98th percentile utilization of 192 kbps, he or she would correctly provision either a 256 or 384 kbps circuit. This would result in a 16% to 25% average and 50% to 75% 98th percentile usage respectively. Based on our analysis, approximately 10% of the circuits were under-provisioned (>50%-60% average utilization), 65% were properly provisioned and 25% were over provisioned (<40% 98th percentile utilization). Given the short duration of the ground campaign and the four fold traffic increase in SIPRNET traffic in less than one month, the results indicate that CENTCOM planners did an excellent job estimating demand and provisioning dedicated bandwidth based SIPRNET networks for OIF.

Table 5 – Percentage of Circuits Exceeding 98th Percentile Utilization Thresholds

Network	98th Percentile Utilization			
	30%	40%	50%	60%
SIPRNET OIF	76%	66%	56%	49%
NIPRNET OIF	79%	71%	62%	55%
NIPRNET	79%	67%	56%	51%

STAT MUX GAIN

The wide variation in circuit utilization described above suggests that fixed provisioning of bandwidth results in an inefficient use of the resource. This inefficiency can be substantially improved with a bandwidth-on-demand network architecture. Satellite BoD networks are able to allocated bandwidth on a real-time basis to meet the dynamic needs of each circuit. The typical architecture provides reserved bandwidth that is always available with no latency beyond the standard geostationary satellite latency. The architecture also provides guaranteed dynamic bandwidth that is available with approximately 500 ms start-up latency. This bandwidth remains allocated to the circuit as long as it is needed. Finally, the network provides best-effort bandwidth on an as-available basis.

In general, a bandwidth-on-demand architecture will provide better resource utilization and improved performance if the underlying network traffic has two characteristics. First, it must be bursty on a node-by-node basis. This was clearly true for the DoD traffic monitored for this paper, as shown by the previous analysis. Second, the aggregate traffic cannot be bursty. In other words, if all nodes require bandwidth at the same time, BoD will be ineffective since there will not be adequate bandwidth to support all user needs. In an effort to evaluate the effectiveness of BoD, the authors aggregated CENTCOM NIPRNET and SIPRNET traffic during Operation Iraqi Freedom and determined the total traffic variation during the course of the operation. Figure 5 shows the mean, 75^{th}, 90^{th} and 98^{th} percentile aggregate demand relative to the peak demand during the operation. This data was generated by adding circuit bandwidth usage for each CENTCOM circuit in each 5 minute sample period. The cumulative distribution of the aggregate bandwidth was computed for each 5 minute period. The data clearly shows the 24 x 7 nature of operations with vary little variation during the course of the day. In addition, the data indicates that the peak aggregate demand during the entire Operation Iraqi Freedom was 45% to 50% higher than the busy hour average demand. In other words, a bandwidth-on-demand network can operate at an average load of 65% to 70% and still have adequate reserves to support the peak demands. Figure 6 shows the aggregate bandwidth CDF relative to average bandwidth. The 98^{th} percentile demand could be met with 46% more bandwidth than the 24 x 7 average bandwidth. The highest demand 5 minute period during OIF (99.99^{th} percentile) required 75% more bandwidth than the monthly average. This compares quite well with the 300% to 400% provisioning bandwidth required using dedicated point-to-point circuits.

The results of this monitoring analysis agree extremely well with simulated network traffic behavior. In fact, analysis indicates that as the number of circuits and the amount of bandwidth increases, the difference between the peak and busy hour traffic will decrease; thus allowing the bandwidth-on-demand network to operate at even higher average load levels. Another advantage of the BoD architecture is that the 10% to 15% under-provisioned circuits would no longer suffer from severely degraded performance since the BoD architecture would dynamically allocate more resources to these nodes when needed.

The collected data suggests that while military operations may be highly correlated, real-time network traffic across a major theater of operations is not. To test the validity of this hypothesis, we performed a correlation analysis for each pair of CENTCOM ports. A total of over 8000 correlation analyses were performed. While the ingress and egress traffic for a given node is highly correlated, we found very few pairs of circuits that were correlated. Table 6 shows the percent of port pairs with various correlation coefficient ranges. The data indicates that less than 1% of circuit pairs were highly correlated (+/-50%), while 8% had moderate correlation (+/-20% to 50%), and 91% were uncorrelated. This is exactly the traffic behavior needed to maximize resource utilization for a BoD architecture.

Table 6 – CENTCOM SIPRNET Port Correlation Statistics during OIF

Correlation Coefficient	Percent Port Pairs
<-50%	0.1%
-50% to -20%	2.3%
-20% to -10%	6%
-10% to 10%	73.7%
10% to 20%	11.0%
20% to 50%	6.2%
>50%	0.7%

Our analysis conservatively estimates that a bandwidth-on-demand SATCOM architecture could reduce transponder capacity requirements by a factor of three and provide adequate headroom to meet peak network demands. The preceding STATMUX gain analysis of DoD NIPRNET and SIPRNET traffic was extremely conservative in estimating gains for BoD architectures for a number of reasons. First, Tier 0/1 routers aggregate traffic from lower level networks thus producing less burtsy traffic and higher circuit utilizations than the networks below them. In the future, these Tier 2/3, and even COTM, networks will connect directly to SATCOM networks, which should increase the STATMUX gain that can be realized for these nodes. Second, we eliminated any circuits with utilization below 1% to insure that the utilization was not skewed lower by backup or unused circuits. Third, we based on estimate on busy hour traffic utilization not 24 x 7 utilization which was slightly lower. The result is a conservative estimate of circuit usage that produced a conservative statistical multiplexing gain estimate.

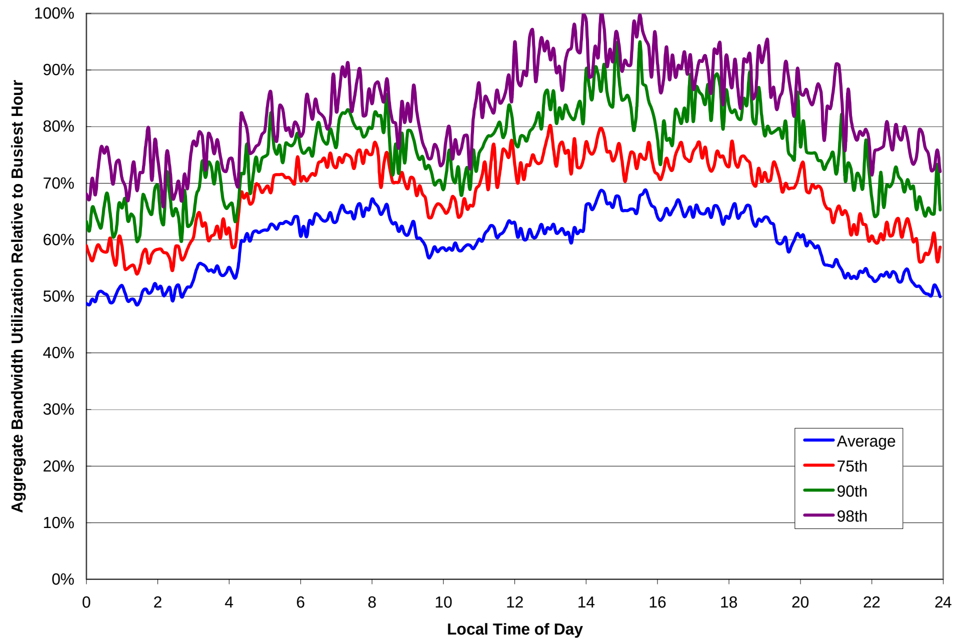

Figure 5 – CENTCOM SIPRNET Time of Day Usage

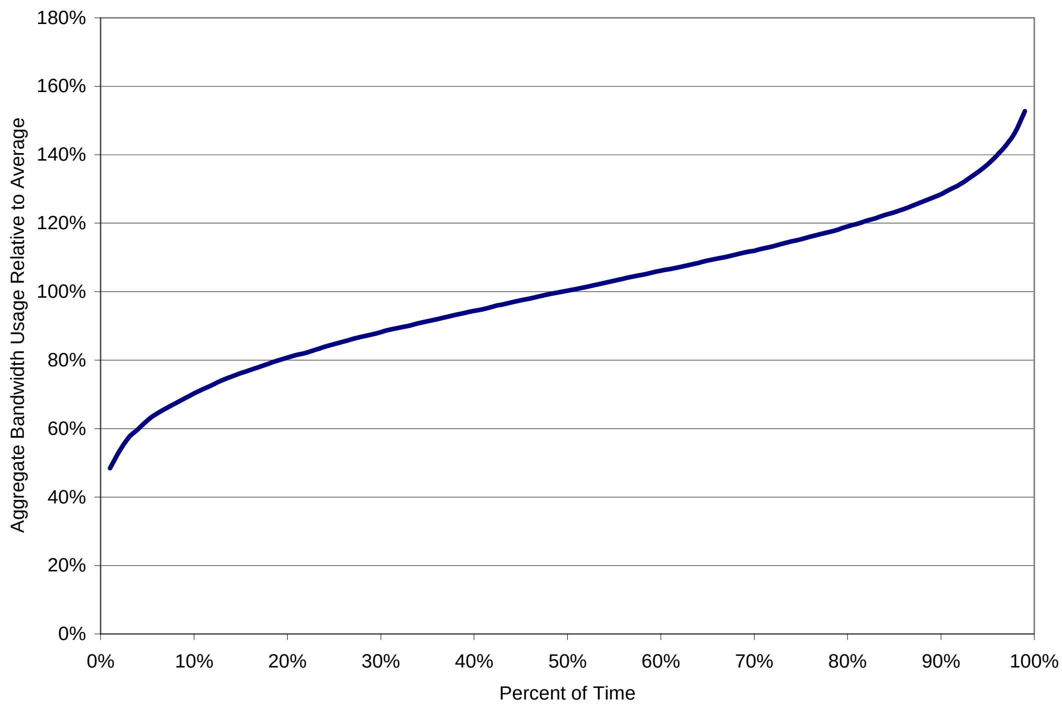

Figure 6 – CENTCOM SIPRNET Usage CDF

CONCLUSIONS

The following conclusions can be drawn from the DoD network measurements. First, on average DoD circuits are under-utilized even during combat operations. Typical utilization during major conflict ranged between 20% and 30% while the utilization during normal operations was between 10% and 20%. In addition, no difference was seen between intra and inter (to CONUS and EUCOM) theater circuit usage. These results are not dramatically different from the 15% usage seen in commercial circuit usage measurements from a previous study. Circuit utilization varied dramatically across an AOR. Closer examination of CENTCOM NIPRNET and SIPRNET circuit usage showed 20% of the circuits operating below 10% average utilization, while 10% of the circuits were operating above 70% average utilization. Second, DoD data traffic is inherently bursty for all circuit speeds. The near-peak circuit demand is typically 3 to 10 times the average busy hour demand. Third, DoD network planner did an excellent job provisioning circuits. We found that circuits were correctly provisioned more than two thirds of the time based on bandwidth and the technology that was available. Analysis of the full 14 month monitoring period indicated that network planners were able to anticipate bandwidth growth and provision additional capacity to maintain fairly constant network utilization even in the face of dramatically increasing demand. Finally, the aggregate circuit bandwidth demand is quite stable even during major combat operations. We found the peak aggregate demand exceeded the OIF average demand by 40% to 50%. The traffic on the vast majority of circuits was uncorrelated and the overall behavior was ideally suited for a transponded or on-board routed bandwidth-on-demand SATCOM architecture. We conservatively estimate that a BoD architecture could reduce SATCOM bandwidth requirements by two thirds as compared to SPCP, FDMA and dedicated bandwidth TDMA architectures.