

Considerations for a Government VoIP Network

John Grass, Director of Engineering, net.com

MAY 4, 2004

Considerations for a Government VoIP Network

Some of the world's largest communication networks are deployed within the various government agencies. Like Tier 1 carriers, they provide voice and data services to numerous users over a wide geographic area.

While government agencies are not driven by revenue and competition, they are interested in data and voice convergence for real cost reductions as well as supporting enhanced functionality for end users. Convergence allows for only one network to run both voice and data networks, removing the burden of having two separate networks. In addition, converged networks may provide enhanced services such as integrated multimedia messaging including voice, video, email, and fax. But introducing convergence into government networks is not the same as in a carrier network. Strong reliance on satellite links and the support of secure voice are two examples of these differences.

BANDWIDTH EFFICIENCY

Over the years, bandwidth-efficient government TDM networks have evolved. While the voice compression capabilities have long been a touted benefit of VoIP, compression in TDM networks over satellite has been the norm in government networks for many years. Hence, a VoIP call with compression should not be compared to a PCM TDM call, but against a compressed TDM call. The compression algorithms in VoIP networks are essentially the same as those available to TDM networks.

Moreover, VoIP calls add a significant amount of IP header overhead on top of the compressed call. As an example, a call compressed with the G.723.1 codec at 6.4 Kbps would generate a data payload of 24 bytes and an IP header of 26 bytes. Compressed TDM calls on the other hand, do not require extra bandwidth beyond the compressed voice.

SILENCE SUPPRESSION

There is, however, one technique that does allow the bandwidth used by VoIP calls to be significantly reduced. This technique is silence suppression. With silence suppression, no VoIP packets need to be transmitted when there is silence. Normally, at least one party is silent while the other party is talking. A call can be made with an immediate bandwidth savings of 50%.

Studies show that on average 62% of a standard voice call is silent. The technique of silence suppression does not apply to TDM calls, as the bandwidth for a TDM call is completely allocated through the network at the beginning of the call. That bandwidth is reserved and constant throughout the call and cannot be shared with other calls regardless of whether or not voice information is sent (i.e., silence suppression is used).

The net result of increasing bandwidth with the IP header overhead with the reduction gained through silence suppression, brings bandwidth usage by VoIP calls equivalent to that of compressed TDM calls.

For VoIP technology then to really make an impact on the bandwidth usage as compared to TDM calls, the large IP header tax needs to be solved.

FRAME PACKING TO REDUCE IP OVERHEAD

One very effective technique to reducing, and in some cases nearly eliminating the IP overhead bandwidth in VoIP networks, is Frame Packing. Frame Packing works as follows – instead of creating an IP packet for every frame of compressed voice data, multiple frames of voice are loaded into the same IP packet.

The net result of increasing bandwidth with the IP header overhead with the reduction gained through silence suppression, brings bandwidth usage by VoIP calls equivalent to that of compressed TDM calls.

There are two dimensions on how to add additional voice frames to the same IP packet. In one dimension, multiple voice frames are added from the same voice call. So instead of putting in one 30-ms frame of compressed voice (e.g., G.723.1 at 6.4 Kbps) into a packet, four 30-ms frames of compressed voice can be put into the same packet. As a result, the IP overhead is effectively divided by factor of four. This is a significant gain.

This technique does have a physical limitation in that only so many frames from the same call can be put into the same packet. Too many frames would have a noticeable delay in getting the voice data to the other end of the network. Users can only tolerate so much delay in their calls, otherwise the call is unintelligible.

A dramatic increase in the number of voice frames that are sent in the same IP packet is achieved by putting in voice frames from different calls that are simultaneously occurring. This means that if there are 10 simultaneous calls, rather than sending 10 separate IP packets for each given frame length of time, one packet is sent with all 10 frames of data. This concept can extend to about 60 calls sharing the same packet. As a consequence, the IP overhead divided by 60 approaches a near negligible amount of bandwidth (i.e., 26 bytes IP overhead divided by 60 gives < 0.5 bytes per call, or < 2% overhead).

A variable to using frame packing is the interval in which packets are sent. For example, if packets are sent every 60 ms and the codec frame size is 30 ms, then two frames from the same call are sent in each packet. A clear advantage is shown using frame packing as compared to not using frame packing (i.e., RTP used in SIP and H.323).

With the IP overhead burden resolved, the bandwidth savings from the silence suppression is really able to shine. The net bandwidth used by the VoIP calls becomes roughly half the bandwidth used by compressed TDM calls. When trying to setup networks where bandwidth is precious, a 50% savings in bandwidth for voice traffic is warmly welcomed!

In order to benefit from the services and functionality developed in the commercial markets, government agencies are best directed to adopt the commercial standards in VoIP, such as H.323 and SIP. These two signaling protocols, however, do not offer the frame packing technology that is key in meeting the bandwidth. As a result, government agencies must look for solutions that can offer advanced VoIP techniques to get the added bandwidth savings for optimized satellite links. Optimization provides cost savings on leased public links and improved network usage on private links.

FRAME PACKING TO REDUCE TOTAL VOICE PACKETS

Putting voice over data networks is a wonderful concept due to the enormous infrastructure and operations savings of running one as opposed to two networks. However, the characteristics of voice and data are not the same. Voice has particular requirements that must be met as well as particular characteristics that must be accommodated by the network.

One particular characteristic of VoIP traffic is the amazingly small size and number of packets that can get generated for a call. While data connections typically involve packets at or near maximum MTU sizes of 1500 bytes, one VoIP call can generate 33 mini-packets of only 50 bytes every second! Consider

hundreds of calls from the same location, and the network now enjoys a flood of small packets. Provisioning for such capacity in the data network can be costly, as sufficiently capable routers that can handle the enormous volume of small packets must be deployed.

Government agencies do have constrained budgets, so the idea of spending considerable dollars to beef up existing data networks to handle the addition of voice traffic must have a significant return. Fortunately, one can achieve significant value without requiring a major upgrade to the data network. As was discussed earlier, frame packet techniques exist whereas voice frames from multiple calls can be packed into the same IP packet. Not only is there the bandwidth savings, there is also the considerable reduction in packets sent out into the network. Consider 50 calls originating from the same location, rather than 1,650 mini-packets of 50 bytes per second, 33 packets of 1,226 bytes could be sent. No over-running the data network after all!

QoS TECHNIQUES FOR VoIP

The voice packets are extremely time-sensitive and consequently must be given priority on the network over bursty data traffic.

The Quality of Service (QoS) given to the voice packets can start right from the source of the voice packets (i.e., VoIP gateway). DiffServ is a standard that was created to enable differentiated services in the network. This is achieved by specifying how priority flags with the IP header can be set. For voice, packets can be flagged as express forwarding, from which the network can be provisioned to give these packets priority over non real-time data packets.

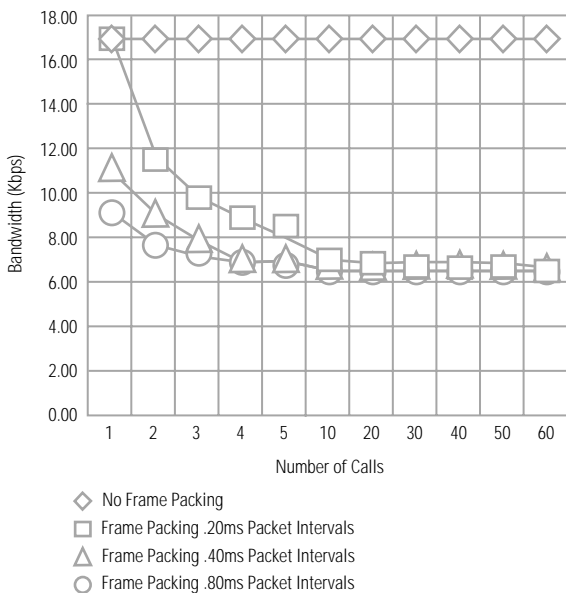


Figure 1: Illustrates the average bandwidth used per call with and without frame packing.

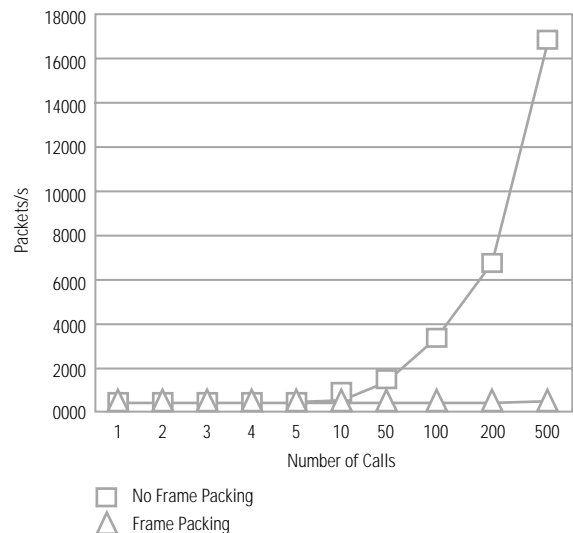


Figure 2: Illustrates the significant number of packets generated without frame packing.

SUPPORTING STU/FNBBDT IN A VoIP NETWORK

Many government networks must support secure communications. There are two main technologies used for secure voice calls. Secure Telephone Unit (STU) is a U.S. federal standard that has been around for many years. STU-III is the latest version of the standard used by U.S. agencies, while STU-IIB is used by NATO.

The basic methodology used by STU enabling secure phone calls is when making a call secure, the voice call is switched to a modem call. Over that modem call, the voice is transmitted as data. In transmitting the voice as data from the phone, advanced encryption techniques can be employed. The voice is encrypted right at the source (the phone). The network need not have any knowledge of the encryption techniques, and only sees the call as a modem call carrying data.

The newer standard for secure voice communications is FNBDT (Future Narrow-Band Data Transmission). From the network point of view, FNBDT is somewhat similar to a STU call. From an unsecure voice call, a modem call is established between the endpoints. Over this modem call, the voice information is exchanged as encrypted data. The modems used by STU and FNBDT are different, as well as the encryption techniques.

In addition to the utilizing secure phone techniques, the entire data stream produced from the secure devices can be bulk encrypted along with other data streams by specialized data encryption devices.

The main consideration of secure call support is that from the network point of view, it is really a special modem call that needs to be transported. While it sounds simple enough, there are large consequences to that fact. All the wonderful bandwidth savings that were discussed earlier that came from voice compression now get thrown into the bit bucket!

MODEM RELAY KEY TO SUPPORTING STU/FNBBDT IN VoIP NETWORKS

While speech compression through the use of codecs is wonderful for bandwidth savings on voice calls, it doesn't work on modem calls. Speech compression works by specifically using the human speech physiology to reduce the amount of bandwidth required to represent the speech signal. Modem signals are not constrained by the human physiology, and hence speech compression algorithms can't be applied to compress the signal bandwidth. As a result, VoIP gateways typically use PCM (64 Kbps) to transfer modem signals. In some instances, attempts to cut this down to 32 Kbps are done by using differential waveform coding (i.e., ADPCM). But the results aren't that great, as the capability of the modem transfer rate gets reduced.

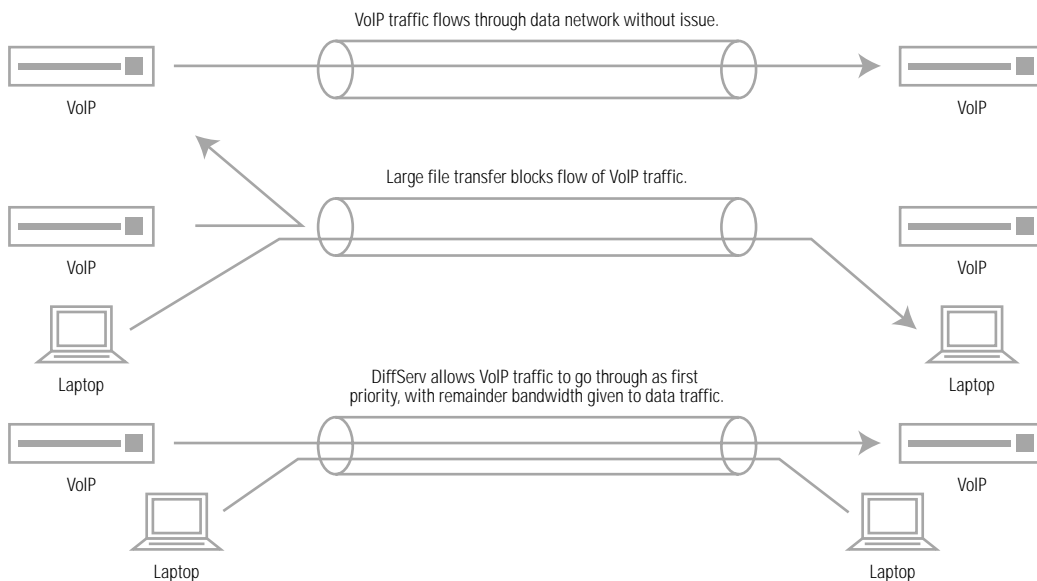


Figure 3: Importance of VoIP prioritization across a data network.

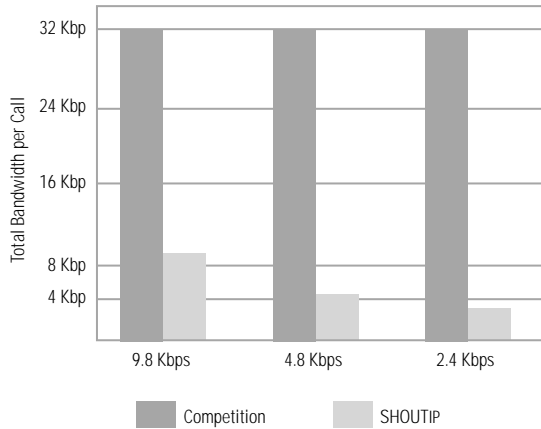


Figure 4: The graph shows advantages of net.com's modem relay transport for STU calls.

NOW THAT I'VE HAD IT, I WANT IT BACK

It is somewhat frustrating that after planning the network to take advantage of the significant bandwidth savings from VoIP speech compression, silence suppression, and frame packing, to lose the benefits when secure calls are introduced into the network.

The answer lies in not using waveform coding of the secure phone modem signal, but rather to terminate the modem signal at the network ingress side. From there, the actual modulated data is recovered from the signal, and only that information is transferred through the network. Secure calls operate at 14.4, 9.6, 4.8, and 2.4 Kbps. So by only sending the demodulated information, the bandwidth requirements are effectively dropped back down to that of the compressed voice. Don't forget to make sure that the demodulated modem information enjoys the same frame packing benefits right along side voice data frames within a packet.

The technology used to demodulate secure calls at the network originating side and only transmitting the demodulated data across the network is called a Secure Call Relay. When planning a VoIP network within a government network, equipment should be evaluated on its ability to offer a Secure Call Relay for both STU and FNBDT. Further, there are hybrid secure phones that offer both STU and FNBDT options on the same device.

FNBDT SPECIFIC BANDWIDTH SAVINGS

For FNBDT Secure Call Relay, asynchronous data is run over the modems. This implies that part of the bandwidth used is for the transmission of the start and stop bits. So an additional 20% of the bandwidth used by the start and stop bits need not be transmitted if the Secure Call Relay is sophisticated enough to properly strip off the bits at the originating side, and correctly reinsert them at the terminating side.

JITTER CONCERNS WITH MODEM RELAY

Jitter in a VoIP network is a pretty common situation dealt with by VoIP gateway equipment. A simple buffer can smooth out many of the packet's inconsistent arrivals while adding a modest amount of delay. Packet behavior beyond the capabilities of the jitter buffer impact the quality of the voice call with poor speech extrapolations or silence gaps. But generally it does not prevent the call from continuing.

For secure voice, the issue becomes more perilous. Since secure calls are a modem call, gaps in the signal sent to the modem are generally fatal. Modems work by tracking the change of the signal received and mapping the changes to data representations. Missing or delaying the signal causes the termination side to become out of sync from the originating side. Then all the data representations constructed from the signal are completely incorrect. The modems quickly realize this and lose sync, returning the secure call back to non-modem, non-secure mode.

A sophisticated VoIP gateway will be able to maintain the modem synchronization between the two sides of the call by being secure modem aware. Any gaps in modem signal are gracefully handled so that both sides maintain sync despite the loss of data. Loss of data on a secure voice call may impact the speech quality, although FNBDT has error correction itself, so voice degradation is minimized. The call, more importantly, remains connected and secure. Dropping out of secure mode requires an annoying amount of time to reestablish the secure call.

ERROR CORRECT TO MAINTAIN SECURE VOICE CALLS

For voice calls, listeners notice when more than a certain amount of speech is missing and due to the time-sensitive nature of voice calls, it's nearly impossible to resend lost packets in a timely manner.

A small amount of tardiness of voice packets is effectively dealt with by jitter buffers. Almost all VoIP equipment has jitter buffers to ensure smooth playing out of voice packets. In a reasonably configured network, only a modest extra amount of delay need be added from the jitter buffer to the call.

For superior call performance, error correction techniques are used to enable immediate correction of corrupt or lost packets. These techniques enable the VoIP equipment to reconstruct lost or corrupt packets immediately without requiring a retransmission from the originating side.

THE IMPACT OF PACKET LOSS ON SECURE VOICE CALLS

Applying error correction techniques to secure calls can be quite beneficial, as the modem's connection can be better maintained in networks with large packet loss.

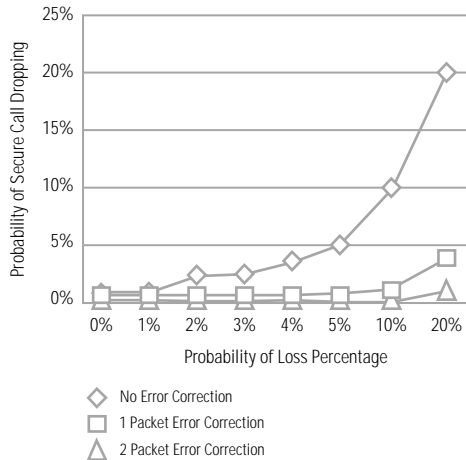


Figure 5: Probability of dropping a secure call with network packet loss.

The graph above shows the probability of a secure call being dropped as a result of different amount of packet loss in the network. As you can see, error correction techniques can greatly improve the reliability of the secure call. The trade-off to adding error correction functionality is the additional bandwidth required on the link to support it. So enabling error correction depends on the quality of the IP network.

THE IMPACT OF SATELLITE DELAYS IN SECURE VoIP NETWORKS

Voice, being real-time data, is very sensitive to delay. Satellite hops introduce 250 ms one-way delay in a network. Two satellite hops, which are not that unusual in some government networks, give a whopping 500 ms delay in one direction, or a full second of round-trip delay.

The issue with long delay is in the network. The voice payload is delayed as well as there is a chance that the secure modems lose sync. The modems expect certain aspects of the modem negotiation and training to happen in a timely fashion. Long delays from satellites compounded by VoIP jitter buffers delay, can make modem success rather elusive.

As a result, advanced secure modem relay that can tolerate long satellite delays in VoIP gateways is a strong requirement in government networks.

BRI INTERFACES FOR SATELLITE TERMINAL EQUIPMENT

Satellites are used within many government voice networks. Therefore, a consideration when specifying the VoIP equipment for locations connected via satellite is to ensure correct connectivity to satellite ground station equipment. As an example, on the commonly used INMARSAT terminal equipment, a BRI data interface is used for the link between the satellite transmitter and the attached network equipment. As such, unnecessary boxes can be avoided if a BRI data interface is available on the VoIP gateway.

SUPPORT OF FIELD OFFICES

Government network providers sometimes have their version of SOHOs (Small Office, Home Office) at the end of a satellite link. What this means is that there is a small number of people requiring both voice and data services. The challenge is to provide the best possible service for both the data and the voice over the small bandwidth satellite connection. Direct control of the voice and data streams is required to manage the precious bandwidth of the link. Satellite bandwidth should only be tapped when required by voice or data services, or otherwise remain unused. Moreover, priority of the voice traffic over the data traffic must be given. An effective VoIP solution would be able to manage both the data traffic and the voice traffic on the satellite link.

The diagram on the right shows an example of a government network where two remote sites are operating secure voice and data services via satellite. The hub site has larger capacity and connects to a PBX.

EQUIPMENT SIZE

Equipment space is always an important consideration for network operators. Power considerations can also play a factor in some locations. For effective network manageability, the less boxes the better. In some locations, sometimes there just isn't enough room for that extra box. So any deployment must consider the number and size of boxes required at a site to enable VoIP. The more requirements that can be met by a single box – such as VoIP capabilities, analog phone interfaces, satellite systems interfaces, frame packing abilities, call control, and voice QoS – the better the solution.

LEGACY SUPPORT

Government networks have long histories, and as a result, have considerable legacy voice equipment in their networks. Since equipment is not being thrown-out, support of this equipment through interconnection is mandatory by any new network gear such as VoIP gateways. Legacy equipment such as PBXs often use older CAS protocols as an example. In addition, the support of such existing applications such as Multi-Level Precedence and Preemption (MLPP) is required. MLPP is a key feature that allows for identified calls to take priority over others.

Government agencies providing networking services are just like any carrier, and are constantly evaluating new technologies that can provide better or more efficient services to their customers. VoIP can provide both new services as well as opportunities for cost reduction. Hence, any new VoIP network enables these possibilities by supporting the latest standards in VoIP protocols such as H.323 and SIP.

MANAGEMENT

The most important aspect of maintaining a large complex network is strong network management capabilities. This has long been a focus of government

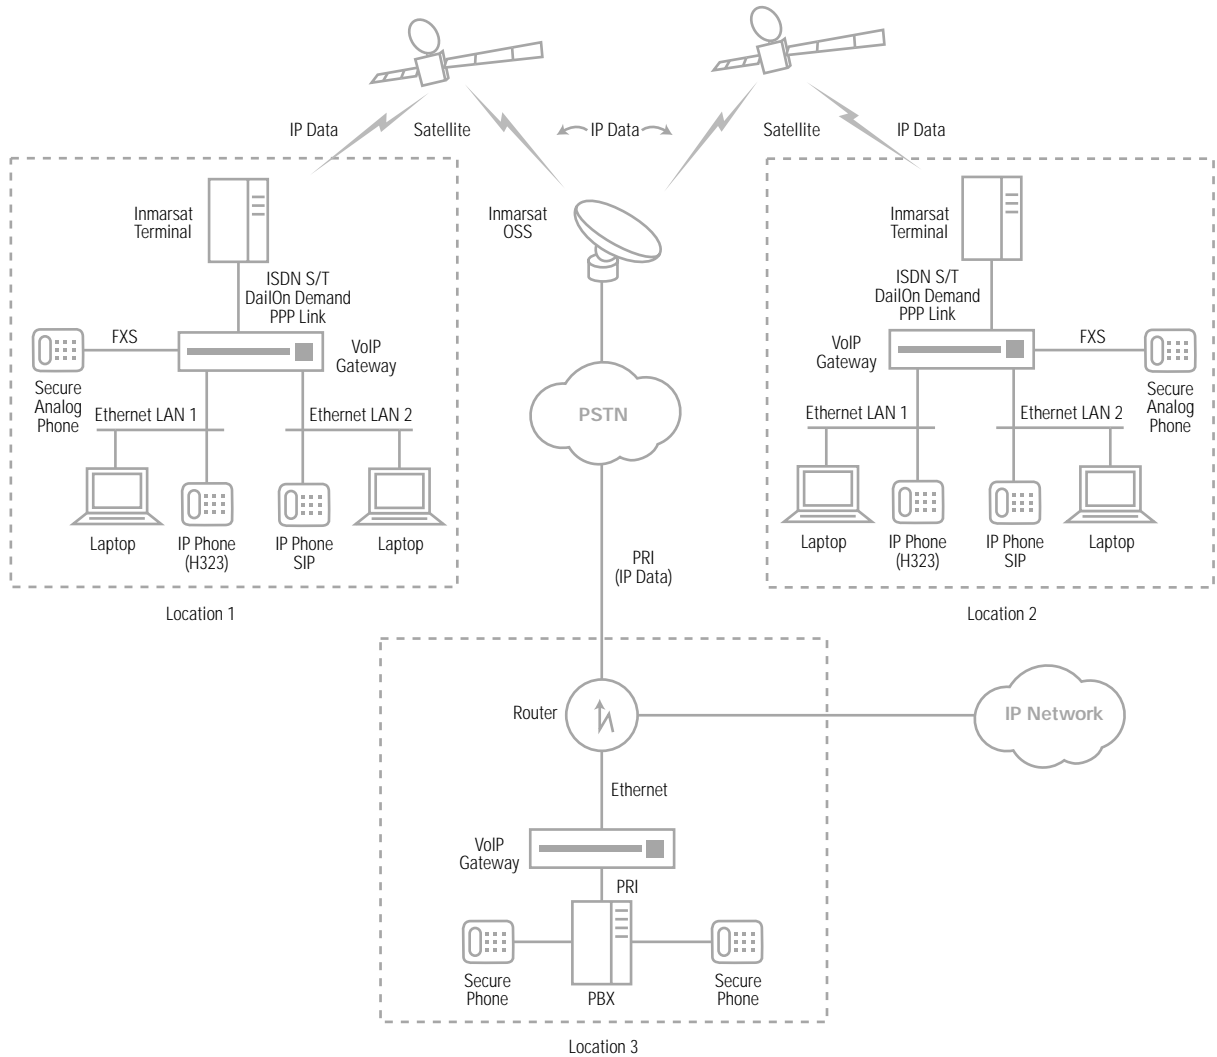


Figure 6: Government network example.

agencies when looking to deploying new equipment. They need to be able to integrate the new equipment with the old equipment and have overall network visibility. In addition, easy-to-use configuration and monitoring tools are very important in reducing operational costs.

SUMMARY

VoIP offers an exciting evolution for government voice networks. Numerous benefits can be realized by the deployment of this new technology. Many similarities between commercial carrier networks and government networks exist. There are, however, some special considerations that government agencies must consider when deploying VoIP.

Bandwidth efficiencies, secure voice, QoS, configuration and monitoring tools as well as a compact solution are some of the key features required in government networks. The key is choosing the right equipment that can address all of these requirements, and not just a limited subset. Silence suppression, framepacking, secure call relay, jitter buffers, QoS, and routing support significantly enhance a basic VoIP offering for government network applications.

An all-in-one solution reduces the amount of equipment to manage as well as offers a smaller footprint solution. The correct VoIP solution can provide government networks both the ability to offer enhanced services as well as reduce operational expenses.



Corporate Headquarters

6900 Paseo Padre Parkway
Fremont, CA 94555 U.S.A.

T 510.713.7300

F 510.574.4000

E info@net.com

www.net.com

N.E.T. Federal

21660 Ridgetop Circle, Suite 100
Dulles, VA 20166, U.S.A.

T 703.948.1800

F 703.948.1850

E net_federal@net.com

Some features listed in the specifications are under development.

© 2004 Network Equipment Technologies, Inc., d.b.a. net.com. Promina, SCREAM, SHOUTIP and IDNX are registered trademarks, and SHOUTbuilder, SHOUTgate, SHOUTscript, SHOUTvue, SHOUTwatch, net.com, and the net.com logo are trademarks of Network Equipment Technologies, Inc., and its subsidiary, N.E.T. Federal, Inc., d.b.a. net.com. All other trademarks and registered trademarks are the sole property of their respective companies. All rights reserved.

This document does not create any express or implied warranty by net.com or about its products or services. net.com's sole warranty is contained in the written product warranty for each product. The end-user documentation shipped with net.com's products constitutes the sole specifications referred to in the product warranty. The customer is solely responsible for verifying the suitability of net.com's products for use in its network. Specifications are subject to change without notice.

SH-WP-0604-600m2