



Application Note

Voice over IP Applications in the LinkStar VSAT System

Introduction

The rise of the Internet and IP as a standard, common form of communication has made IP the protocol of choice for various applications over both wireline and wireless communications media. One of those applications that has grown in popularity over recent years is Voice over IP (VoIP). In response to the quickly evolving trend towards VoIP, ViaSat has enhanced the LinkStar product to address and overcome the specific VoIP challenges unique to satellite. This effort has resulted in a highly competitive LinkStar Voice over IP application, particularly when incorporated with ViaSat's recommended voice appliances.

This paper discusses VoIP considerations in the LinkStar system and illustrates how LinkStar provides high quality, cost-effective Voice over IP, thus minimizing total cost of ownership (both equipment and bandwidth costs). Application examples are provided at the end of the document, in conjunction with helpful tips and information.

VoIP Considerations for Satellite

Due to the expense of satellite space segment, satellite service providers are necessarily more conscious about bandwidth utilization than their terrestrial counterparts. The bandwidth required for transporting the actual Voice over IP must be minimized to provide a viable business model.

Voice over IP, by nature, is very inefficient as the payloads are typically very small (~20 bytes), and have IP headers (IP/UDP/RTP – 40 bytes) that are usually bigger than the payloads themselves, in this case producing VoIP packets with an efficiency of 33%. For terrestrial networks, this relative inefficiency may not be a concern since bandwidth is virtually unlimited, however, as mentioned earlier, the expense of satellite space segment makes it imperative that any viable satellite VoIP application seek to minimize overhead.

Equally important in VoIP applications (satellite or terrestrial) is the ability to prioritize VoIP traffic ahead of less time-sensitive traffic, such as email and file transfers, to avoid the delay and jitter that can render a voice application unusable. It is therefore critical to have Quality of Service (QoS) capability within the network to identify, prioritize, and perform rate controls.

Finally, service providers typically consider a VoIP signaling protocol when implementing VoIP applications over satellite. This is much less critical than bandwidth usage or QoS, but is a common issue about which service providers ask for guidance. The various signaling protocols differ in the size and quantity of call setup messages

transmitted back and forth over the satellite. Since each message going across the satellite and back results in ~540 milliseconds of satellite delay, the desire is to use protocols that minimize this handshaking.

LinkStar's innovative and efficient VoIP application takes into consideration the three elements discussed next in this document:

- ♦ VoIP Overhead (Packet and Packing Efficiencies)
- ♦ IP Quality of Service (IP QoS) and TDMA Access
- ♦ VoIP Signaling Protocol

VoIP Overhead (Packet and Packing Efficiencies)

Inbound

Uncompressed Header

In order to better understand the inherent efficiency issues with VoIP, we need to look at the details of what happens when packetized voice is encapsulated into IP payloads. In this example, we consider what happens to a G.723 (5.3 kbps) voice packet over an IP stream. We will assume a default sample size of 20 bytes.

Standard VoIP with 5.3 kbps G.723 would look something like this:

The coder will output a 20-byte payload every 30ms, as shown in the calculation below:

$$(20 \text{ bytes} \times 8 \text{ bits/byte}) / (30 \times 10^{-3} \text{ seconds}) = 5,333 \text{ bits / sec (or 5.3 kbps)}$$

Therefore, we have a 20-byte payload, which then has a 40-byte header added to it (Exhibit 1).

IP Header = 20 bytes

RTP Header = 12 bytes

UDP Header = 8 bytes




Exhibit 1 – VoIP Packet Structure

(Note: The 14-byte Ethernet header and 4-byte CRC are not mentioned because they are stripped away at layer 2.)

Now, we have a 20-byte payload + 40-byte header, which produces a 60-byte payload every 30 ms. The result is that the 5.3 kbps voice call now requires 16 kbps over IP!

Proof of calculation is shown below:

$$(60 \text{ bytes} \times 8 \text{ bits/byte}) / (30 \times 10^{-3} \text{ seconds}) = 16,000 \text{ bits / sec (or 16 kbps)}$$

Compressed Header

In the satellite business, we cannot afford to use 16 kbps of satellite throughput to place (1) 5.3 kbps voice call. For this reason, the LinkStar system employs header compression to reduce overhead inefficiency.

Since VoIP user traffic does not have a well-known port number, it is not possible to compress the header based upon a commonly recognized port number. To solve this issue, the LinkStar system does the following: it looks at each UDP packet coming in and based upon other heuristics looks for a valid RTP header. The LinkStar remote (RCST) then computes whether a reduction of the 40-byte header to a 3-7 byte header uses fewer LinkStar PDUs (protocol data units, or segments) on the return link. If so, the RCST compresses the header from 40 bytes to approximately 3-7 bytes, resulting in a greater than 80% reduction in the VoIP header!

Additionally, in order to improve LinkStar packing efficiency, ViaSat's implementation uses larger voice payload sizes. For example, this G.723 voice call can be coded using a 40-byte payload every 60 msec.

This is illustrated below:

$$(40 \text{ bytes} \times 8 \text{ bits/byte}) / (60 \times 10^{-3} \text{ seconds}) = 5,333 \text{ bits / sec (or 5.3 kbps)}$$

After LinkStar header compression and the addition of the compressed header (7 bytes as worst case), we obtain:

$$(47 \text{ bytes} \times 8 \text{ bits/byte}) / (60 \times 10^{-3} \text{ seconds}) = 6,267 \text{ bits / sec (or 6.3 kbps)}$$

LinkStar actually allocates slightly more than 8 kbps of bandwidth for this 6.3 kbps voice call. However, since the voice packet rate is only once every 60 msec (versus 45.5 msec TDMA rate), there are times when other traffic (such as other data) can be interleaved with the voice. This provides maximum usage of the TDMA payloads.

See Appendix A with a summary of the various coding techniques tested and the inbound throughput required.

Outbound

LinkStar does not perform header compression on the outbound because the open standard DVB-S encapsulator must be able to recognize the VoIP packets. However, inefficiency is far less an issue on outbound since the outbound stream uses a very efficient TDM stream that also benefits from voice activity detection and statistical multiplexing on its multiple outbound voice streams. This voice activity detection is estimated to reduce bandwidth usage in the range of 30-40%.

IP Quality of Service (QoS) and TDMA Access

Inbound IP QoS

IP QoS can be implemented using techniques such as Integrated Services (IntServ (e.g., RSVP)), Differentiated Services (DiffServ), and MPLS. LinkStar implements IP QoS using the Differentiated Services (DiffServ) implementation based on RFCs 2474, 2597, and 2598.

Using DiffServ, the LinkStar system effectively supports latency-sensitive services such as Voice over IP by providing a comprehensive QoS that identifies and tags high priority traffic, prioritizes that traffic ahead of other lower priority traffic, and polices/shapes the flows as needed.

First, LinkStar can be configured to identify or tag traffic in order to assign it to a class of service. By enabling IP QoS in LinkStar, the remote terminal will either tag an IP packet that meets a pre-defined flow profile (according to source IP address, destination IP address, source port #, destination port #, or protocol), or it will recognize an IP packet that has already had its codepoint set by an external device. The DiffServ model defines 6 classes of service using different codepoints. These classes of service (as implemented in LinkStar) from highest to lowest priority are: EF (Expedited Forwarding), AF1, AF2, AF3, AF4 (Assured Forwarding 1, 2, 3, 4), and BE (Best Effort). Once the traffic has been identified and tagged appropriately, the traffic must be separated according to its class.

LinkStar supports 1 queue according to the DiffServ model for each class of service listed above. These 6 queues are applicable for the TDMA side (exit queue) of the LinkStar remote terminal and are designed to avoid the serialization delay encountered when using just 1 queue over slower WAN links.

For voice over IP, the codepoint can be set to EF (expedited forwarding) to provide top priority. This means that the LinkStar terminal will ALWAYS send a segment waiting in the EF queue before any other segment waiting in lower priority queues (AF1-4 and BE). This is independent of the amount of bandwidth the remote terminal has. It simply means that if there is a burst available on the remote terminal's TDMA link and there is a segment in the EF queue, that segment will always go onto that burst before segments in other queues. The remaining queues will be emptied in priority after the EF queue. This helps to guarantee that a voice call is not disrupted in the presence of other bursty, bandwidth-intensive traffic, such as email with attachments or file transfer, regardless of the voice call's packet size and arrival time.

The final step in assuring that voice over IP traffic runs smoothly is to utilize rate control. Depending upon the network design, the LinkStar remote terminal will have some amount of bandwidth available to it, whether CIR, CIR-Reallocation, BOD, or some combination. Therefore, the operator can define some rate policies to control the rate of flows for the VoIP application, as well as define policies for other applications desiring to share the bandwidth available to the remote terminal. Within the LinkStar system, the operator can define policies (maximum rate of flows and burst lengths in conjunction with the codepoint setting), to determine the action taken on flows in which bandwidth limits are exceeded. This allows the operator to divide up and prioritize bandwidth allocations as needed to support the applications from the remote.

Inbound TDMA Access

Before leaving the inbound IP quality of service discussion, we will review the effect of the inbound TDMA access scheme employed. This is not “IP QoS,” but overall QoS that we need to discuss.

LinkStar can support VoIP services using several methods. One method is through the use of a best-effort access scheme to TDMA bandwidth. This is more commonly termed bandwidth on demand (BoD). BoD is a best-effort bandwidth allocation scheme whereby traffic bursts are dynamically allocated to remotes (RCSTs) only when they request bandwidth by sending a bandwidth request message to the Hub. In this way, the limited returnlink bandwidth is allocated only when needed, and only as much as needed, thereby allowing a common bandwidth pool to be shared among many users. This is the most efficient (and cost effective) method and is common for applications, such as rural telephony, where large numbers of subscribers share a small amount of available bandwidth.

There are some VoIP applications (such as enterprise telephony), however, where a higher degree of service is warranted. For these applications, LinkStar has the capability to configure committed information rates (CIR) that provide dedicated bandwidth that is always reserved for that remote site. This enables the service provider to guarantee certain bandwidths and performance for each site in the network.

In order to support CIR efficiently, ViaSat’s advanced bandwidth allocation scheme also supports a feature called CIR-Reallocation (CIR-R). CIR-Reallocation provides a dedicated amount of bandwidth to a remote site only if the site needs bandwidth. If the site is idle (and thus requiring no bandwidth), the CIR-R bandwidth allocated to that site is given to the BoD pool available for any site to use. When the idle site becomes active with traffic, the bandwidth is pre-empted from the BoD pool and given back to the site that was formerly idle. This provides maximum flexibility with highly efficient bandwidth usage.

The net result is that as the voice traffic load increases, high quality voice services are maintained and the service provider can guarantee the quantity of voice calls determined during network sizing.

Outbound

On the outbound link, the LinkStar QoS manager identifies and classifies traffic, provides prioritization, and performs rate control giving minimum/maximum bandwidths in similar ways that the inbound link implements quality of service. This is accomplished through the use of an optional QoS manager.

VoIP Signaling Protocols

VoIP Signaling type is not nearly as crucial as the packet efficiencies or QoS, but worthy of brief discussion since so many clients inquire about the various VoIP signaling protocols available. The most common protocols in use today are H.323, SIP, and MGCP. (Also worth mentioning is SCCP, a Cisco proprietary protocol that is fairly common to Cisco VoIP networks.)

H.323

H.323 is the most common protocol and has been used for quite some time. This may be the safest protocol to use for those who would like to use a widely available configuration. However, H.323 was not developed with the WAN link in mind.

Therefore, there is a bit more signaling for call establishment/teardown that takes place with this protocol in comparison with a protocol like SIP. However, in many cases, this may be indistinguishable to the user (especially for more recent versions that minimize call setup delay). There have been security notices issued that indicate that there may be some security liabilities using H.323 in some Cisco devices, but workarounds and upgrades are available to address these issues. (See Cisco advisory entitled “Vulnerabilities in H.323 Message Processing,” initial public release 13 January 2004.) For such reasons, operators may want to consider alternative protocols.

SIP (Session Initiation Protocol)

The Internet Engineering Task Force (IETF) developed the SIP standard for multimedia calls over IP. ASCII-based SIP works in client/server relationships, as well as in peer-to-peer relationships. SIP uses requests and responses to establish, maintain, and terminate calls (or sessions) between two or more end points.

SIP is one of the newer protocols, and as such may not be supported by all equipment manufacturers. If your VoIP equipment supports this protocol, then it is worthy of consideration as it is relatively “light” in terms of signaling traffic for call establishment/teardown. This protocol does not use a lot of commands to establish calls. Basically, there are only 4 exchanges for the setup (“invite,” “try,” “connect,” and “ok”) for a simple call scenario. When using SIP, coding type, voice activity detection, etc., are generally defined at the endpoints, not the call manager. This simplifies the call setup and minimizes the number of messages exchanged over the satellite.

MGCP (Media Gateway Control Protocol)

Media Gateway Control Protocol (MGCP) is a plain-text protocol used by call-control devices to manage IP telephony gateways. MGCP (defined under [RFC 2705](#)) is a master/slave protocol that allows a call-control device (such as CallManager) to take control of a specific port on a gateway.

MGCP (Media Gateway Control Protocol) has been described as a very “light” protocol in that it does not utilize large call setup messages, but does a fair number of exchanges. It is easy to interoperate with and has mainly been used with cable phones and other analog endpoints. There are a couple of different ways to initiate calls. Like SIP, this protocol was actually designed with some bandwidth constraints in mind. MGCP tends to be used by devices that are fairly unsophisticated in nature, and consequently is common to low-cost analog telephone adapters at remote sites. With this said, there are a fair number of exchange messages used during setup to define various parameters since the end devices rely on the master for all configuration and instruction.

SCCP (Skinny Client Control Protocol)

SCCP uses Cisco-proprietary messages to communicate between IP devices and Cisco CallManager. SCCP easily coexists in a multiple protocol environment. The Cisco IP Phone is an example of a device that registers and communicates with Cisco CallManager as an SCCP client. Cisco ATA186/188 telephone adapters are another example of Cisco devices that can use SCCP. (Note: Cisco devices can also use other standard protocols as described above.) During registration, a Cisco IP phone receives its line and all other configurations from Cisco CallManager. Once it registers, it is notified of new incoming calls and can make outgoing calls. The SCCP protocol is used for VoIP call signaling and enhanced features such as Message Waiting Indication (MWI).

H.323, SIP, MGCP, & SCCP have all been implemented successfully over the LinkStar network. In some cases, the protocols do not wait for responses before proceeding to the

next step, so call setup times can vary some between the different protocols and their configurations.

Application Examples

LinkStar has been used in several different applications with VoIP. The optimum selections of VoIP application equipment and/or configuration settings vary with the application. A couple examples of LinkStar with VoIP are included below.

Examples

Rural Telephony

Exhibit 2 below displays LinkStar as used in a rural telephony type application with a large number of remote sites attempting to access the PSTN (or possibly other remote sites via double-hop connection.) This application is typically characterized by a need for low-cost analog telephone adapters providing FXS ports mainly to telephones and possibly fax machines. There are typically a large number of voice channels to be aggregated, so it is very important to have a large gateway at the hub that can handle large call volume. Furthermore, the hub is typically characterized by a very robust gatekeeper and call manager to support the required call routing, call setups/teardowns, call processing, calling features, etc. Some applications require the ability to provide pre-paid or post paid calling. Pre-paid calling and other calling features can be provided by purchasing those additional options.

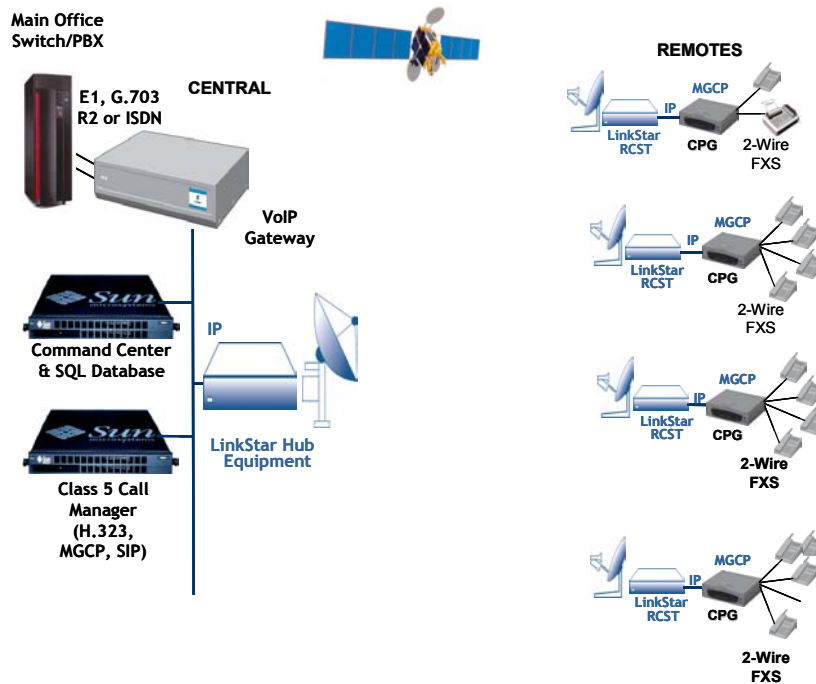


Exhibit 2 – Rural Telephony Application

Enterprise Telephony Application

Another VoIP application found in LinkStar is the enterprise application where the remotes may connect to phones, fax machines, PBXs, or central office lines.

This application is typically characterized by a need for low-cost analog telephone adapters providing FXS ports and higher end gateways that can connect to PBXs to provide E&M or E1 interfaces. There are typically more channels per site than rural telephony, but fewer sites in total. This application still requires a fairly robust gateway at the hub that can handle the voice traffic. Furthermore, the hub is typically characterized by a very robust gatekeeper and call manager to support the required call routing, call setups/teardowns, call processing, calling features, etc. Enterprise applications are more likely to require additional calling features such as caller ID and three-way calling, among others.

It is also important that the VoIP equipment can provide T.38 fax relay so that fax calls will generate ~25kbps IP traffic load as opposed to ~80kbps IP traffic load. (T.38 is an ITU-T standard for fax relay.) Some inexpensive customer premise gateways (CPGs) do not support T.38 fax relay and will not properly respond to the tones (not voice as expected). Therefore, the codec will switch from its compressed mode to PCM (64k) mode and generate approximately 80 kbps of IP traffic load. However, when using fax relay, the fax/modem tones are demodulated at the CPG and then passed as data. CPGs that support this T.38 fax relay will pass this fax real-time producing approximately 25 kbps of IP traffic load. ViaSat can provide VoIP appliances configured to support T.38 fax relay for Group 3 (G3) faxes and super G3 faxes.

In Exhibit 3 below, you will notice that MGCP is used for the lower cost VoIP gateways and SIP used for the higher end gateways providing connections to PBXs.

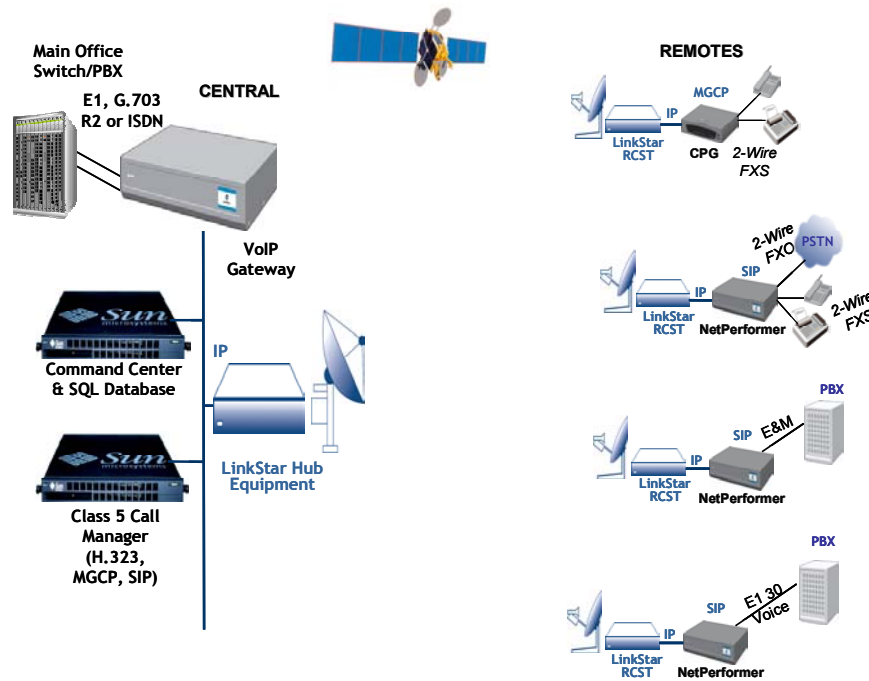


Exhibit 3 – Enterprise Telephony Application

Conclusion

ViaSat’s cutting edge LinkStar platform has been developed with features such as VoIP in mind. This foresight has led to a platform that addresses packet inefficiencies, IP QoS, and overall QoS as it relates to access schemes, all of which empower service

providers with the potential for success in using this application. These characteristics make LinkStar the platform of choice for VoIP.

Appendix A

VoIP Bandwidth Utilization in LinkStar

Manufacturer / Coding Rate	Voice Compression / Payload Size / Qty of Calls	Duration for Sample (msec)	Maximum Quantity of PDU's / Frame
Cisco			
5.3 kbps	G.723 (40 bytes) – 1 call	60	1
	G.723 (40 bytes) – 2 calls	60	2
8 kbps	G.729 (40 bytes) – 1 call	40	2
	G.729 (40 bytes) – 2 calls	40	3
Verso/Clarent			
5.3 kbps	G.723 (40 bytes) – 1 call	60	1
	G.723 (40 bytes) – 2 calls	60	2
8 kbps	G.729 (40 bytes) – 1 call	40	2
	G.729 (40 bytes) – 2 calls	40	3

Note: Each LinkStar PDU provides slightly more than 8 kbps of inbound throughput.

Appendix B

The table below provides a partial list of manufacturers whose equipment has been implemented over the LinkStar system.

Manufacturer	Model	Interface Provided	Protocols Used	Call Manager/ Gatekeeper Used?	Location Used (Hub or Remote)
Cisco	3725 Router	E1 – R2 signaling (China #1)	SCCP	Yes	Hub
	ATA 186	FXS	SCCP	Yes	Remote
	1751 Router	FXS/FXO	H.323 (no gatekeeper)	No	Hub and Remotes (demo purposes)
Clarent (Verso)	BHG 1000	E1- R2 Signaling	MGCP (only)	Yes	Hub
	Netperformer	FXS/FXO/E&M/E1	SIP	No	Hub and Remote (demo purposes)
	VG101/VG201	FXS	MGCP (only) SIP to come soon	Yes	Remote
Audio Codes	MP-104 / MP-108	FXS	H.323 (no gatekeeper)	No	Remote

Appendix C

Recommendations / Helpful Hints

1. Instead of cRTP, LinkStar uses a proprietary implementation of header compression designed specifically to optimize VoIP efficiency over LinkStar. Therefore, when using VoIP over LinkStar, do not turn on cRTP if using Cisco equipment (or other devices that support cRTP). Instead, enable header compression on LinkStar and let it perform the necessary header compression.
2. Use G.723 (5.3k) voice compression and set the voice payload size to 40 bytes if possible. This is the most efficient VoIP solution for LinkStar and its header compression is optimized around this standard. The second best VoIP solution for LinkStar is G.729 (8k) voice compression – the payload size should be set for 40 bytes here as well. These payload sizes are configurable on most VoIP equipment.
3. If your network will carry fax, then VoIP appliances available through ViaSat are highly recommended. They support T.38 fax, which will produce about 25kbps of IP traffic load, as compared to 80 kbps of IP traffic load for non-T.38 fax support. These versions available through ViaSat also run special software that allows them to detect not only G3 faxes, but also super G3 tones as well and pass them as T.38 fax. Other units will see the tones and pass this using standard PCM producing ~80 kbps over IP. (As of the initial date of this document creation, Cisco ATA 186/188 does not support T.38 fax – Cisco indicates that it is in the product roadmap.)
4. For significant quantities of remotes (20 or more), you will want to use a call manager so that routing tables can be managed from one location and easily updated. Otherwise, with each change, you will have to change 20 routing tables, one for each device in network. ViaSat engineering assistance is available to help you select an optimum VoIP solution.
5. Use LinkStar inbound QoS to make sure that VoIP traffic is prioritized and does not get impacted by data services such as FTP. See LinkStar documentation or ViaSat system engineering for assistance on setting up QoS in the network.

ViaSat, Inc.

Corporate Headquarters

6155 El Camino Real
Carlsbad, CA 92009-1699
Phone: (760) 476-2200
Fax: (760) 929-3941

Clarksburg, MD

22300 Comsat Drive
Clarksburg, MD 20871
Phone: (301) 428-4500
Fax: (301) 428-4700

Norcross, GA

4356 Communications Drive
Norcross, GA 30093-2999
Phone: (678) 924-2400
Fax: (678) 924-2480

Publication Information

Revision Number	Date Released	Comments
001	November 9, 2004	REQ-009672

© 2004 ViaSat, Incorporated. All rights reserved.

LinkStar™ is a trademark of ViaSat, Inc.

ViaSat Proprietary—Information in this document is subject to change without notice. This document is proprietary to ViaSat, Inc., and shall be protected by a receiving party in accordance with the terms of its contracts and agreements with ViaSat, Inc., covering LinkStar™ products.

No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without the express written permission of ViaSat, Inc.