Application of IP Multicasting to the NASA Communications Command and Telemetry Ground Network

M. Chris Spinolo
NASA Goddard Space Flight Center

ABSTRACT

The NASA Communications (Nascom) Division has been directed to deploy Internet Protocol (IP) based technology for the ground segments of all present and future spaceflight telemetry networks. The Nascom network supports all NASA spaceflight telemetry, command and status requirements, from sounding rockets and balloons to the Hubble Space Telescope and the Space Shuttle. This paper discusses the challenges of transitioning a 35 year old, custom engineered, worldwide legacy telemetry network to IP, and the resulting, new NASA IP Operational Network for ground transport of spacecraft telemetry and command.

KEY WORDS

IP, IP Multicast, NASA, Ground Network.

INTRODUCTION

The NASA strategic goals of dramatically reducing costs of routine spaceflight operations and positioning these operations so that they are more readily outsourced, resulted in the directive to eliminate all proprietary and custom data communications systems, using IP based technology instead. This directive was immediately applied to the mission critical telemetry systems as well as the administrative and support networks. The Nascom program managers also added the additional requirement that there should be no impact to the existing applications or additional costs for the customers, other than their acceptance testing of the new network transport.

The re-engineering effort was organized into two large parts:

1) the network-application interface (where the network meets the customer)
2) the network transport infrastructure

Each part then had three major phases:

1) legacy evaluation
2) evaluation and application of related IP technology
3) deployment

The discussion of this paper will follow this outline, but the highlight of this report is the new application of IP multicast technology to the mission critical telemetry network of NASA.

THE NETWORK-APPLICATION INTERFACE

Legacy evaluation

The legacy network interface is based on the familiar clock and data bit stream, captured in the RS-449/422 physical and electrical standards. In order to switch the telemetry in the legacy network, it is formatted into blocks with source and destination addressing, and other information useful to spacecraft and instrument control. The format of these blocks was determined more than 20 years ago, and yet it still serves its purpose well enough today. The focus of this re-engineering effort is the actual switching and network systems that transport these blocks.

Traditionally, this interface has been implemented with specialized hardware, but as general data processing systems have become more powerful, more of the data formatting, synchronization and data reduction is being done in software. However, because this legacy goes back as much as 35 years, there are still very many, now archaic, processing subsystems in use. Many of the legacy telemetry blocking and de-blocking subsystems are dependant on the synchronous and continuous nature of a clock and data serial bit stream. This re-engineering effort was not authorized nor funded to update these front-end systems, and was in fact directed not to impact these customer interfaces in any way. Taking advantage of the fact that the legacy network backbone uses domestic satellite carrier services, buffers in the new network application interface were designed to be large enough to smooth variations in the inter-packet arrivals of an IP transport. In this way, there is always enough bits in the buffer to provide a continuous bit stream from the serial interface. However the limited delay of a satellite hop, 250 msec, is still considered a constraint that must be addressed in any IP network. In this case it is addressed in the operational concepts discussed below in the Network Transport Infrastructure.

Technology Application

The legacy network infrastructure has been built on two, interconnected, schemes of transport: 1) Point to point circuits connected directly to “Message” Switching Systems (MSS), and 2) Point to point circuits connected to large programmable Multiplexer-

DeMultiplexer (MDM) systems, whose aggregate channel is connected by a satellite broadcast system. The MDM is so integrated into the Space Network at White Sands Ground Terminal, (ground station for the Tracking Data Relay Satellite System, or TDRSS), and Space Shuttle operations at Johnson Space Center, that it was decided to be a much lower risk to make those MDMs IP capable. One of the MDMs at GSFC, and the MSS were replaced outright by the conversion devices discussed next.

The concept is to provide a “conversion device” that interfaced to the legacy serial interface on one side and a nominal IP multicast network on the other. Because there are several hundred serial interfaces in the legacy network, an inexpensive solution (several serial interfaces for less than $10,000) is desirable. There are 2 candidates, both based on PC platforms and special programmable serial interface cards. The first was an in house effort using the Linux OS and a NASA serial interface card, and the second is a product marketed by Avtec called a Programmable Telemetry Processor (PTP) which uses the OS2 or Windows NT OS’s and Avtec programmable serial card interface cards. After several application software patches, both solutions have been deployed, and the PTP is the officially sanctioned solution since, it is commercially available and supported. Another company, General Data Products, has demonstrated their own conversion device, with up to 15 serial channels capacity, to be compatible with the NASA IP transition network.

NETWORK TRANSPORT INFRASTRUCTURE

Legacy evaluation

The legacy transport network is a serial switched and multiplexed network, providing each project a scheduled, dedicated channel for each required stream. The specialized switches, using mainframe based communication processors, also performed a lot of data replication so that a single, real time stream could be received by several destinations at once. The multiplexed data was transmitted over satellite carrier services, which also performed an inherent data replication for multiple destinations.

There is no distinction of data types, and so all data, whether spacecraft commands, attitude telemetry or day old spacecraft recorder playback data is treated equally, as critical to the life and success of the mission. There are very little transport services provided by the legacy protocol, primarily providing error detection and sequence detection. These characteristics have significant physical layer performance implications, which are discussed later.

IP technology application

IP multicasting, a recent development in the Internet, is another product of the Internet research community. Its application to distance learning, conferencing, and communications integration (i.e. data, voice and video) has spawned considerable commercial interest and development, to the point that there is currently a worldwide “Internet within the Internet” known as the Multicast backBone or Mbone, which relies on this technology, and is quickly becoming a commercially viable entity, such as the World Wide Web.

The application of IP multicasting technology is implied by the capability of the legacy switching systems to copy telemetry blocks, and the use of satellite carrier services, which inherently duplicate telemetry by the number of earth stations receiving the downlink in the MDM system . Effectively applied, IP multicasting can also utilize wide area bandwidth very efficiently. IP multicasting is considered mature since several router manufacturers support it, the Internet has sponsored it for several years, and it is available in several low cost host platforms, including the candidate conversion devices. The primary drawback is that IP multicasting supports only datagram (UDP) transport services. The option of using a proprietary or special application as a reliable transport protocol was ruled out, for increased risk, lack of schedule to test and parameterize, and it was considered anathema to the directive to use standard generally available protocols. The issues of relying on UDP as the IP transport protocol are addressed in the deployment section.

Telemetry stream group management

The Internet Group Management Protocol (IGMP) handles the set up of a telemetry stream by associating a telemetry source with 1 or more destinations. IGMP defines a protocol between hosts (the conversion devices) and the IP routers. A typical telemetry stream could have the telemetry source as the ground station in contact with the satellite, and the telemetry receivers are the mission control centers (with online and backup receivers), science operations centers, flight dynamics and navigation facilities, data production and other facilities. During launch and early orbit, the number of telemetry receivers often multiplies when specialists, investigators and dignitaries request their own real time displays for these exciting and critical times. The dynamic nature of the group memberships lends itself well to IGMP which reacts on demand. The scheduling systems and protocols which support the legacy network are considered burdensome and costly. The use of IGMP for these purposes eliminates the scheduling function and its associated systems.

Telemetry Stream Routing

A routing protocol is required so that the IP routers know how to route the multicast data based on the availability of network resources, such as the routers themselves and the circuits and subnets connected to them. There are now several flavors of multicast routing protocols to choose from, and they each have distinct behaviors. The behavior NASA requires consists of the following:

1) No data loss for stream setup (first block gets through, important for commands) or routing overhead
2) Deterministic behavior, diagnosable faults
3) Fault tolerant, auto recovery from physical layer failures
4) Efficient use of bandwidth, some telemetry streams are more than 2Mbps
5) Data driven, unscheduled route control (not manually configured any particular telemetry stream)
6) Adequate monitoring capability

There are 3 contenders currently available for IP multicast routing protocols: DVMRP, PIM, and MOSPF. All three protocols rely on IGMP to request that a particular multicast data flow be routed toward a particular network destination. Only the determining characteristics of these three protocols as they would be used in the NASA real time telemetry network are discussed below. These characteristics, or lack of them, were demonstrated in a network laboratory. The actual operation of the protocols is well documented by the references in the bibliography.

DVMRP

The Distance Vector Multicast Routing Protocol (DVMRP) is the defacto standard in the Internet's Multicast Backbone (Mbone). This protocol was developed to allow routers to handle thousands of multicast groups with the least impact on router resources. DVMRP routes multicast data based on explicit rejection notices, in other words, it makes sure everybody gets all the data and processes rejections to keep the data from going to distinct network destinations. This is also known as a dense mode protocol, since where there is a high density of receivers, this approach may work efficiently. DVMRP implements this behavior by flooding the multicast data to all network interfaces that are running DVMRP, which requires little processing by the router. The NASA telemetry network has many "tail" sites which receive only 10's kbps telemetry over narrowband circuits. The flooding behavior of DVMRP would deny these tails sites of their required telemetry at various times for various periods. The use of static controls to restrict this behavior is not acceptable, as they become very numerous for hundreds of telemetry streams and hundreds of destinations. DVMRP is not deemed usable.

PIM, Sparse Mode

The Protocol Independent Multicast protocol was developed by Cisco routers to address many of the issues with DVMRP. It is the leading protocol of use in the Mbone, as a majority of the Mbone connected routers run PIM. PIM sparse mode routes data based on explicit requests. Multicast data is not sent to a network destination unless a receiver at that destination explicitly requests it. While PIM runs primarily in a single vendor's router (Cisco), it is the predominant protocol in the Internet Mbone. PIM requires the configuration of a rendezvous point (RP), which coordinates all multicast flows between each source and all its destinations. In our testing and verification in the lab, the PIM protocol has been very liable to topology failures. The ability to reroute around failed circuits and routers was unpredictable at best and non functional at worst. PIM is also a developing protocol, with new versions expected to be released in the future. The prospect of re-evaluating and regression testing future versions of the protocol are very unattractive prospects for our purposes. Finally, PIM is only supported in Cisco routers, effectively sole sourcing our network infrastructure. PIM was ruled out as a primary choice of multicast routing protocols but highly desired over DVMRP, since under nominal conditions it functioned quite well.

MOSPF

The Multicast Open, Shortest Path First protocol is OSPF with multicast extensions. OSPF is the standard unicast IP protocol in the Internet. MOSPF extensions were developed, implemented and deployed by an Internet Engineering Task Force (IETF) sanctioned working group several years ago, whose work has been documented in the references at the end of this paper. MOSPF performed extremely well in our evaluation tests, behaving in a deterministic manner in both normal conditions and failure conditions. MOSPF is supported by the 2nd and 3rd largest router vendors, offering a competitive commercial source for the needed IP network infrastructure. The quick, deterministic failure recovery and competitive sources of MOSPF were the 2 leading reasons for its selection as the NASA IP Operation Network multicast routing protocol. There have been allocations the MOSPF is too burdensome on the routers to deploy in a large scale. This has proven to be untrue in our experience.

Deployment

The use of the PTP outsourced the production, maintenance and support of the legacy serial interface. For instance, any patches for the PTP have been provided by the vendor under standard warranty. I will say only briefly, there has been a trade-off between having industry assume development and sustaining engineering costs, and schedule control that an in house project would have. NASA choose to accept that risk.

The MDM IP interface remains a custom built, in house sustained, effort. The IP Multicast interface for the MDM replaces the aggregate channel that connects the MDMs over the satellite broadcast transport. The serial interface ports are mapped to IP multicast addresses in the multiplexer and the IP multicast addresses are mapped to serial port addresses in the de-multiplexer. These functions are accomplished in a new processor with a 100BASET Ethernet based IP interface, which are integrated in the MDMs.

The approach is to build the IP infrastructure with adequate bandwidth and monitor the use of bandwidth such that datagrams are not lost in the network due to congestion. The circuits used for wide are transport are provisioned under an existing Government contract that provides premium performance for NASA critical requirements. These are the two primary conditions in the IP infrastructure which allow for the use of UDP transports between applications. We have been aware the NASA receives exceptional bit error performance on its T1s, acquired through its contract mechanisms with its FTS provider (AT&T). NASA contracts for a BER is 1 x 10 –7 and availability of 99.95%, and usually receives better.

It should be noted that MOSPF does have a dense mode behavior in some circumstances. In a non backbone areas with multiple paths to the backbone, multicast traffic is sent to both the primary and backup connections to the backbone. This will result in data flooding the area if the configuration of that area does not take this behavior into account. With proper IP architecture and design this has not been an issue for the NASA network.

A principal reason for the success of the NASA IP Operational Network has been the successful implementation and staffing of a Network Operations Control facility. The NOC relies primarily on well motivated engineer level staff and operators. Its resources are built around Unix workstations running HP Openview and other SNMP tools and applications.

Summary

The conclusion to this matter, for the purposes of this paper, is that IP multicast technology, and MOSPF in particular, has been demonstrated to be mature and ready for real time, critical application. The time has now arrived where telemetry and command networks can be commercially provided and relied on. However, real time critical networks are not yet a commodity, but require careful end to end engineering and analysis.

The most significant area of follow on investigation and analysis is the use of emerging IP multicast transport protocols which offer assured delivery, and recovery of lost packets, as applied to a real time application that requires continuous data.

NOMENCLATURE

COTS	Commercial off the Shelf
COMMGR	Communications Manager
DEMUX	Demultiplexer
DVMRP	Distance Vector Multicast Routing Protocol (standard Internet multicast routing protocol)

GSFC	Goddard Space Flight Center
IGMP	Internet Group Management Protocol
IONET	IP Operational Network
IP	Internet Protocol
JSC	Johnson Space Center
MDM	Multiplexer/Demultiplexer
MOSPF	Multicast OSPF
MSFC	Marshall Space Flight Center
Multicast	when data is sent to multiple destinations simultaneously
MUX	Multiplexer
NASA	National Aeronautics and Space Administration
Nascom	NASA Communications Division
OSPF	Open, Shortest Path First (Internet routing protocol, unicast only)
PIM	Protocol Independent Multicast (Cisco proprietary IP multicast routing protocol)
PTP	Programmable Telemetry Processor
SCD	Small Conversion Device
SNMP	Simple Network Management Protocol
STGT	Second TDRSS Ground Terminal
UDP	User Datagram Protocol
Unicast	when data is sent to a single destination only.
WSGT	White Sands Ground Terminal
WAN	Wide Area Network

BIBLIOGRAPHY

RFC’s

Internet Group Management Protocol, Version 2.
http://www.merit.edu/internet/documents/rfc/rfc2236.txt

Multicast Extensions to OSPF
http://www.merit.edu/internet/documents/rfc/rfc1584.txt

MOSPF: Analysis and Experience
http://www.merit.edu/internet/documents/rfc/rfc1585.txt

OSPF Version 2
http://www.merit.edu/internet/documents/rfc/rfc2328.txt

Protocol Independent Multicast-Sparse Mode Protocol Specification
http://www.merit.edu/internet/documents/rfc/rfc2117.txt

Distance Vector Multicast Routing Protocol
http://www.merit.edu/internet/documents/rfc/rfc1075.txt

RFC Index
http://www.merit.edu/internet/documents/rfc/INDEX.rfc

other internet information:
Internet Assigned Numbers Authority
http://www.iana.org/

Other primary FTP repositories for the RFCs are:

- NIC.DDN.MIL
- WUARCHIVE.WUSTL.EDU
- NISC.JVNC.NET
- FTP.CONCERT.NET
- VENERA.ISI.EDU
- FTP.CONCERT.NET

Textbooks

Kumar, Vinay. Mbone Interactive Multimedia on the Internet. New Riders Publishing, Indianapolis, 1996.

Moy, John. OSPF Anatomy of an Internet Routing Protocol. Addison-Wesley, 1998.
http://www.readmedotdoc.com/0201634724.html

Other resources

The Mbone Information Web
http://www.mbone.com

Mbone archive and Frequently Asked Questions (FAQ)
ftp://isi.edu/mbone

IP MultiCast (IPMC) Initiative
http://www.ipmulticast.com

Reliable Multicast Transport Protocols by GlobalCast Communications.
http://www.gcast.com

NASA IP Operational Network links

[NASA] CIO EXECUTIVE NOTICE 09-95 Network Protocol Standard;
http://www.hq.nasa.gov/office/codea/codeao/notice9.html

IP Transition Project Transition Plan
http://skynet.gsfc.nasa.gov/transition/docs/overview/trans-pln/iptxpln.html

Nascom 4800-Bit Block to Internet Protocol Conversion Description
http://skynet.gsfc.nasa.gov/transition/docs/overview/trans.info.html

Conversion Devices
http://skynet.gsfc.nasa.gov/transition/docs/condev.html

Multiplexer/Demultiplexer (MDM)
http://skynet.gsfc.nasa.gov/transition/docs/mdm.html

NASA Communications Division
http://skynet.gsfc.nasa.gov

Avtec Systems
http://www.avtec.com/sgs