

The Gemini Observatory Multi-continental Network

Gemini Preprint #89

James R. Kennedy
Gemini Observatory
670 N. A'ohoku Place
Hilo, HI 96720

*James R. Kennedy, on behalf of the Gemini IS team

The Gemini Observatory Multi-continental Network

James R. Kennedy^{*}, on behalf of the Gemini IS team
Gemini Observatory^{**}

ABSTRACT

The Gemini Observatory operates two 8-meter IR/optical telescopes: one in Hawaii and the other in Chile. High-speed network connections among all of the mountain telescopes, their sea-level bases, and a support facility in Tucson are essential to their operation, providing video and audio communications, administrative computing systems, remote telescope operation, scientific data management, and many other applications. All the sites have recently been connected via the Abilene network, through collaborations with more than a dozen astronomy facilities near of the various Gemini sites, with Florida International University's AMPATH program, with various providers, and with grant support for the National Science Foundation. While the bandwidth levels required will change over time, Gemini's current objective is a minimum 10 Mbps presence on Internet2 to and from its principal sites. The Gemini North, Gemini South, and Tucson sites are at this level or better. Gemini North has been upgraded to a burst capability up to 155 Mbps to the US Mainland Internet2. Gemini South has burst capability to 10 Mbps, with 6 Mbps guaranteed to the US Mainland Internet2.

Keywords: astronomy, communications, networking

1. INTRODUCTION

Modern astronomy requires access to the full sky, not a view limited to half the universe from only a Northern or a Southern Hemisphere site. The science agencies of seven nations, the US, UK, Canada, Chile, Australia, Argentina, and Brazil formed the Gemini partnership to build and operate two 8-meter infrared/optical telescopes, one in each hemisphere. Between the two, the entire sky is visible at angles suitable for cutting-edge research. The northern telescope is located at nearly 14,000 feet on Mauna Kea on the island of Hawaii. The southern telescope is located at 9,000 feet atop Cerro Pachón in Northern Chile. The two telescopes began science observations in 2000 and 2001 respectively.

The telescopes themselves are amongst the largest in the world, and are uniquely fine-tuned for operation at the infrared (IR) wavelengths accessible at the bottom of the Earth's atmosphere. In recent years the IR spectrum opened up for large-scale observations due to significant technological advances in IR array-detector technology. As a result, today very sensitive solid-state array detectors are available in fairly large formats. Already there are mosaics of CCD chips sensitive to optical radiation in service on Gemini at the 24-Mpixel level. Soon, equally large IR detectors will become available that will allow astronomers for the first time to see and measure the results of physical processes in the Universe that simply cannot be seen in shorter wavelengths with the same clarity and scale as in the optical regime.

The use of these arrays on large telescopes has two significant impacts on the rapid transfer and storage of data. The first is that the arrays are very sensitive and can capture images in minutes or even seconds, that took hours only a few years ago with earlier techniques on smaller telescopes. The second reason is the sheer pixel size of the arrays themselves. For example, consider a 16-Mpixel array producing 16-bit-pixel images at the rate of 90 per hour, with a 10-second readout. The instantaneous data transfer rate from the instrument is 26 Mbps. The rate of accumulation of data is 2.9 GB per hour and over 32 GB per night, and twice that if the data are actually stored in a 32-bit format.

Consequently, the Gemini operations plan for science, engineering, and administration vitally depends on high-bandwidth, low-latency network connections between its principal sites and its science community. At the highest level of priority, these sites include its telescopes and base facilities, its facilities in Tucson, and the data archives.

^{*} jkennedy@gemini.edu, 808 974-2515, Gemini Observatory Northern Operations Center, 670 N. A'ohoku Place, Hilo HI, 96720, USA

^{**} Supported by the Gemini Observatory, which is operated by the Association of Universities for Research in Astronomy, Inc., on behalf of the international Gemini partnership of Argentina, Australia, Brazil, Canada, Chile, the United Kingdom and the United States of America.

The next level includes the National Gemini Offices in the seven Partner countries and any remote observing sites, via Internet2. These are followed by institutional and individual users in the Partner countries, via either Internet2 or the Commodity Internet. Finally, Gemini seeks to be connected to the large pool of interested research and educational institutions, and the general public in the Partner countries and the world, normally via the Commodity Internet. In all these categories, the Observatory is responsible for making itself available on the network, while the clients must fund whatever level of connectivity they desire from their side.

Wright, Fujioka, and Kennedy¹ described the general framework of the Gemini Observatory network structure as of 1999. There it was noted that Gemini North would connect to Internet2's Abilene network (a closed high-performance research Internet), as the result of a \$600,000 grant from the National Science Foundation (NSF), which also enabled the connection of the other Mauna Kea observatories. It was further noted that the NSF had just awarded another grant to Gemini to build an OC3 (155 Mbps) microwave backbone between the Gemini and Cerro Tololo Interamerican Observatory (CTIO) base facilities in La Serena Chile, and their mountain-top telescopes. This two-hop system would go from La Serena to Cerro Tololo, and then to Cerro Pachón to form the city-summit private link for the facilities.

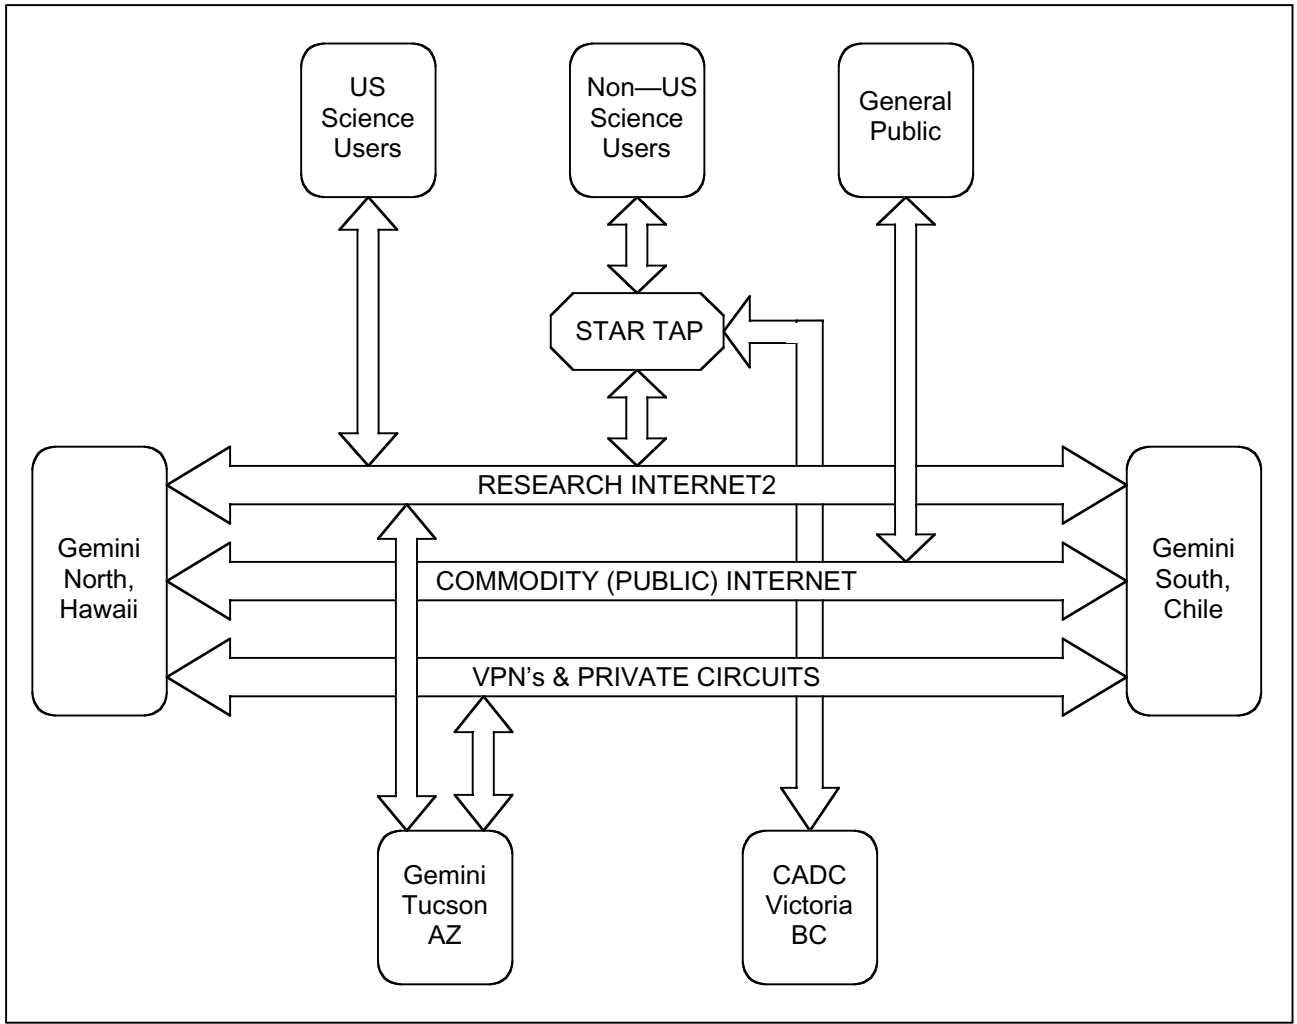


Fig. 1: Simplified diagram of the logical Gemini network between its own internal sites, the archives (CADC, Victoria, Canada), the US and Non-US Partner science users, and the general public worldwide.

2. THE GEMINI NETWORK TODAY

Today the Gemini Observatory facilities consist of two *operational areas* (Gemini North in Hawaii and Gemini South in Chile), a support office in Tucson Arizona, and a third-party archive in British Colombia — the Canadian Astronomy Data Center (CADC). While functionally integral to the Observatory conceptual system, the CADC represents an external entity. The remaining sites are truly internal.

There are multiple logical Internets and other circuits connecting the various Gemini internal networks. These include the Abilene network, the Commodity Internet (the common public Internet), planned Commodity-Internet-transported Gemini virtual private networks (VPNs), and completely private Gemini network circuits.

2.1 Internet2

The principal exchange medium for most Gemini scientific work is the Internet2 Abilene high-performance research network. Access to Internet2 is by membership only, providing two relatively lightly-loaded, cross-connected backbone paths across the US Mainland. Internet2 allows access to more than 300 member and non-member US institutions, including universities, colleges, research centers, observatories of various kinds, and other educational institutions.

Gemini North — Gemini established its Internet2 connection to the facilities in Hawaii in March 2000. Gemini North shares this circuit with the other Big Island observatories and the University of Hawaii. At present the total Internet2 circuit capacity is 155 Mbps.

Gemini Tucson — Gemini's office in Tucson has had Internet2 access for a number of years through another co-located sister organization, the National Optical Astronomy Observatories (NOAO). In turn, this circuit links to Abilene through the University of Arizona. In principle, this shared Internet2 connection is at 2 x 155 Mbps, although the circuit is often heavily loaded with University and other traffic.

Gemini South — The Internet2 connection to the Gemini facilities in Chile began in April 2002. Gemini South shares this circuit with its sister observatory, the Cerro Tololo Interamerican Observatory (CTIO) in La Serena Chile, and its affiliated programs. The circuit is provided by a commercial provider in Chile (ENTEL) from the Gemini/CTIO base in La Serena to the AMPATH² node at the Florida International University in Miami, and thence to Abilene in Atlanta. At the present time, the total Internet2 bandwidth from La Serena to Miami is 10 Mbps, with Gemini enjoying a 6 Mbps PVC within that circuit.

With the completion of this connection, all the internal Gemini sites are available to the science community on Internet2. Moreover, for the first time, the various Gemini sites themselves are now interconnected by a high-performance network.

Science Users — Astronomers in the US can access Gemini directly through Internet2. Internet2 is also connected to an international research network exchange point in Chicago called STAR TAP. This provides peered access to dozens of national and regional high-performance research networks worldwide.

The combination of Internet2 and its connection to STAR TAP provides a means for its science users to access all the Gemini sites. Gemini science users then must provide themselves with the means of connection to the Internet2 backbone or to the international backbones linked to STAR TAP.

2.2 Commodity Internet

Gemini also makes itself available to the general public for information, education, and other related activities. This is accomplished by its connections to the Commodity Internet. Currently, the available bandwidth varies from a burst maximum of 155 Mbps at Gemini North to 4 Mbps sustained at Gemini South.

2.3 Private Circuits

Despite their geographical separation, the Gemini telescopes and their bases and other support facilities represent a single observatory entity. Thus, it is important to have the capability to join the "behind-the-firewall" local-area networks of the local sites into a secure, wide-area network. This permits flexible intercommunication free from the complexities of firewall constraints.

Gemini North and Tucson — For some years Gemini has had a T1 (1.5 Mbps) private circuit connecting its facilities in Hawaii and Tucson. This circuit connects these two facilities seamlessly behind their respective firewalls. Although higher bandwidth is available between the sites on Internet2, this mode does require transiting two firewalls. Staff members at the two sites have the flexibility of choosing the best route to use, based on their performance requirements.

Gemini North and Gemini South — As of August 2002, the private connection between Gemini North and Gemini South has not been completed. In principle, Gemini could establish a *virtual private network* (VPN) between the two operational areas using either Internet2 or the Commodity Internet. However, Internet2 use policy forbids the establishment of VPN's over their circuits. Another option is a leased private circuit, but this is prohibitively expensive.

However, a practical option is the acquisition of additional, Gemini-only, bandwidth between Gemini South and the US Commodity Internet. This is an economically viable approach that could provide a 2 Mbps circuit, and represents the first phase of a plan to complete this needed connection. The current target is to complete such a circuit by the end of calendar 2002.

Internet2 engineers have proposed another possible alternative that may be even more attractive. While their policy precludes a VPN, Internet2 is provisioned with the most recent release of the Internet Protocol (IPv6) middleware. This version has certain secure connection capabilities that may provide the advantages of a VPN, without actually using the VPN protocol. This may represent a phase two approach to connecting North and South behind the firewall.

TABLE I
Gemini Circuits as August 2002

| Point to Point | Bandwidth | Provisioning | Status |
|---------------------------|-------------------------|---------------------|-------------|
| Mauna Kea - Hilo | 45 Mbps | Dedicated fiber | Operational |
| Hilo - Com. Internet | 155 Mbps max | Shared MKO/UH | Operational |
| Hilo - Internet2 | 155 Mbps max | Shared MKO/UH | Operational |
| Cerro Pachón - La Serena | 77 Mbps | Dedicated microwave | Operational |
| La Serena - Com. Internet | 4 Mbps | Shared, CTIO/CARSO | Operational |
| La Serena - Internet2 | 10 Mbps max, 6 Mbps min | Shared, CTIO/CARSO | Operational |
| Tucson - Com. Internet | 155 Mbps | Shared, NOAO/UA | Operational |
| Tucson - Internet2 | 155 Mbps | Shared, NOAO/UA | Operational |
| Archives - Internet1 | TBD | TBD | |
| Archives - Internet2 | TBD | TBD | |

2.4 Operational Areas

Each of the two *operational areas* consists of a sea-level base facility (Hilo and La Serena), and a mountaintop telescope facility (Mauna Kea and Cerro Pachón). The Arizona-based support office is embedded in the NOAO-Tucson base facility and its network infrastructure.

Conceptually, each of the two operational areas has a single border interface device (router) at the sea-level base facility. This border router is a filter that stands between all the traffic on the various Internets and only passes that traffic destined for either the base facility itself or its corresponding telescope facility. The border router then separates the sea-level traffic from the mountain traffic and passes these to those specific sites where the traffic is delivered to separate firewalls protecting each location.

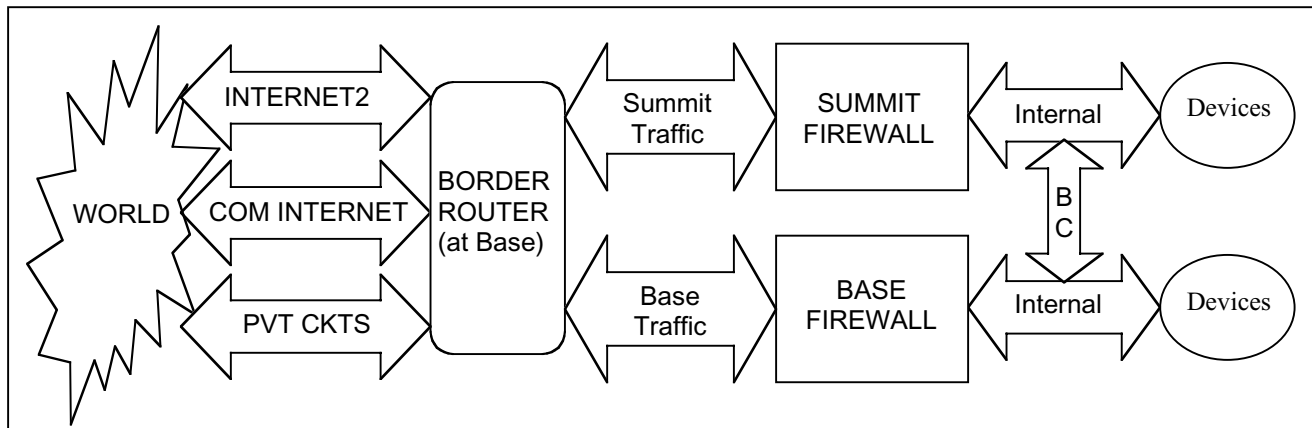


Fig. 2: The conceptual interface between the "outside world" and the base and summit sites of a Gemini *operational area* (i.e. Gemini North in Hawaii, or Gemini South in Chile). Note the Back Channel (BC) permitting direct communication between the Base and Summit Facilities.

Traffic passed inward by the firewalls is sent to the addressed devices by one or more switches to various virtual LANs (vLANs). Within each *operational area* there is a wide bandwidth *back channel* that allows trusted, direct communication between the Base and Summit facilities.

The primary function of the back channel is to connect the Base and Summit control rooms to permit a seamless interface between them during telescope observing operations. It also has the function of separating direct telescope-operations traffic from general Internet traffic. The main impact is that it allows each Gemini telescope to be operated from two locations; virtually identical control rooms are located in both the summit *and* base facilities of each observing site.

Though conceptually identical, the actual implementation of the Fig. 2 configuration is somewhat different between Gemini North and Gemini South. These detailed differences arise from the fact that the Gemini North network stands completely alone, in the sense that it does not share its border router with any other entity, and that it relies on a private DS3 (45 Mbps) fiber circuit for the Base-to-Summit back channel circuit.

By contrast, Gemini South is embedded within the joint Gemini/CTIO network and thus shares its border routers with CTIO and other on-site facilities, and it shares a Base-to-Summit OC3 microwave feed with those same facilities, to provide the Base-to-Summit back channel with a 77 Mbps PVC.

3. FUTURE PLANS

Gemini's bandwidth needs are a moving target driven by the ongoing development of new instruments and detectors, and also by the growing application of network technologies to other issues within the observatory. There are two challenges: 1. predicting the bandwidth needs of the Gemini Observatory and its client astronomers over time, and, 2. predicting the requirement to adapt work practices and Observatory methodologies (which are currently based on a "low bandwidth" Observatory operations paradigm) to make optimal use of the new opportunities enabled by the availability of reliable, high-bandwidth links. Nevertheless, it is clear that there will continue to be a need to accommodate progressively higher bandwidths as these developments unfold.

More specifically, Gemini will be installing one or more solutions to establish a functional equivalent of a behind-the-firewall Ethernet connecting all of its internal sites. The missing link currently is the connection between Gemini North and Gemini South. The Commodity VPN solving this problem will be opened before 2003.

Telepresence is a central functional requirement. Gemini already extensively uses H.323 video conferencing over the Internets to coordinate and manage its day-to-day activities between the sites, and with the various Partner entities, many of which are vitally connected to the science planning and the development of next-generation instruments, as well as direct science collaboration. The next step is to implement a comprehensive telepresence system that allows Observatory-wide "virtual teams" to interact via video, sharing documents and using virtual white-boards. This need is continuing to grow and is having an impact on the evolution of behind-the-firewall facilities.

Another addition in the offing is the integration of facets of the observatory's telephone system with the Internet using Voice over IP (VoIP). With its far flung sites, it would be very convenient if, within the observatory, the telephone sets at all sites appeared simply as extension numbers to any other site.

With four control rooms, one on each summit and base facility of each telescope, in principle a control room in the South could be used to operate the telescope in the North, and vice versa. Already there is a requirement for remote diagnosis and control of engineering functions across the sites, since the Observatory frequently has one specialist per system who (by definition) cannot be in two sites at once. While not yet implemented, this could be a useful capability under a number of circumstances, and the current network bandwidth is sufficient to support such an arrangement.

Both of Gemini's operational areas are quite remote to the vast bulk of the Partnership's science community. As a natural extension of the above concept, there is vision that certain selected sites in the Partner countries would be allowed to establish "remote viewing rooms". These would be centrally located facilities that would resemble one of Gemini's actual control rooms.

This would allow teams of scientists to congregate at a convenient central location, but not require overseas travel, to participate in the actual observational activities remotely. They would be able to see the screens the Gemini operators see, converse with the operators by video conference, and basically do everything *except* control the telescope itself.

The NOAO in Tucson, CTIO in Chile, and the University of Florida all have plans to provide such a facility for their constituents. Experiments in using a rudimentary form of this mode, and even in directly controlling instruments from a remote site, have already been successful.

At the present time, the Gemini Science Operations plan calls for the movement of science data from the telescope to the archives, and the distribution of archival data to the community, on CD and DVD media. However, the availability of higher bandwidth in the future calls into question the economics and reliability of these processes. It might well be less expensive, especially in terms of staff time, if the transfer to the archives took place via the Internet. Notwithstanding, having the option to move archival data to end users via the network certainly could be an attractive feature as well.

It is anticipated that network technology will continue to meet the demands of modern astronomy. But as in the past, this will require the continued support of thoughtfully developed national technology-development efforts such as those ongoing at the NSF³ and other places.

ACKNOWLEDGMENTS

While having some role in its development, the author is largely a reporter of the vision and accomplishments of the many other people on the Gemini networking team over the last several years. These include J. Wright, K. Gillies, T. Fujioka, F. Bull, L. Bass, N. Saavedra, J. Peysson, K. Cornwell, D. Pedreros, D. Melder, C. Aspin, J. Garcia, C. Boyter, N. Bersamina, and at our sister observatory CTIO, T. Ingerson, R. Lambert, and C. Smith.

The Gemini Observatory is operated by the Association of Universities for Research in Astronomy, Inc., under a cooperative agreement with the NSF on behalf of the Gemini partnership: the National Science Foundation (United States), the Particle Physics and Astronomy Research Council (United Kingdom), the National Research Council (Canada), CONICYT (Chile), the Australian Research Council (Australia), CNPq (Brazil), and CONICET (Argentina).

REFERENCES

-
- ¹ J. R. Wright, T. Fujioka, and J. Kennedy, "Network Infrastructure Improvements at Gemini", Gemini Preprint 53, 1999.
 - ² J. Kennedy ed., *Identifying Areas of Collaboration Between the U.S. and the AMPATH Service Area*, The AMPATH Program, Florida International Univ., 2001
 - ³ T. DeFanti and M. Brown eds., *NSF CISE Grand Challenges in e-Science Workshop Report*, Elec. Visualization Lab., Univ. Illinois at Chicago, 2002