

# STScI Bandwidth and the Archive

STUC Presentation

11/12/2009

Carl Johnson

# Prologue

- STScI established an Archive Team to unify all our multi-mission archive services, operations, and resources under a single team.
- One goal for this team is to improve the access and delivery of data to users.
  - ◆ Increasing our external bandwidth was identified as an area that could have an immediate effect.
  - ◆ Started looking to see how much our external connection speed affects data retrieval.
- Network performance issues are no surprise:
  - ◆ Comments in user surveys
  - ◆ Study/Presentation in 2004
    - ◆ Resulted in OPO moving externally.
    - ◆ Goddard increased bandwidth allocation.
      - They had limited this below the physical 100 Mbps.



# Today's Presentation

*"Does the STScI Archive have sufficient connectivity or throughput to support its user community? "*

- Presentation Contains:
  - ◆ Introduction
  - ◆ Current State at STScI
  - ◆ Current Mitigation Activities
  - ◆ Affects of Increasing Network Loads
  - ◆ Closing Comments

# Introduction

## ■ Internet at the border (STScI)

- ◆ External network is 100 Mbps to Goddard that can route to Internet or Internet 2.
- ◆ Internal network has 10 Gbps backbone and is 1 Gbps capable to the desktop.

## ■ What do others have (or are moving to)


- ◆ Canadian Astronomy Data Centre (CADC): 1 Gbps Internet 2 moving to 10 Gbps
- ◆ University of California Campuses: (3) 1 Gbps with 1 connection dedicated for research
  - ◆ Run by Corporation for Education Network Initiatives in California (CENIC)
- ◆ University of California Santa Barbara (UCSB): (6) 1 Gbps moving to 10 Gbps.

## ■ Network Considerations

- ◆ Slowdowns take time to clear out.
  - ◆ Good analogy for networks are highways.
- ◆ Maximum capacity and throughput is difficult to achieve.
- ◆ Understanding peak loads is more important than average loads.
  - ◆ Volume across the network may be greater than actual file size due to packet wrappers, re-transmission, etc.




# Current State



- Screening Router Data showing traffic in and out of the building.
  - ◆ As the timescales increase, the bin sizes also increase.
    - ◆ Peaks begin to get hidden by non-peak periods.
  - ◆ Network charting tools do not preserve the underlying data
    - ◆ Makes doing any follow-up analysis more difficult
- Questions we wanted to answer.
  - ◆ Are there regular peaks in network usage (normal rush hours)?
  - ◆ How do archive requests contribute to network usage? Do they align with the peaks?
  - ◆ What kind of performance are individual external users seeing in their retrievals?
  - ◆ What kind of increases can we support given current network and user trends?

# Current State: Peaks

Average Bandwidth  
8/14/09-9/10/09




- 75-80 Mbps is where performance begins to rapidly degrade.

- Took underlying network data and binned to hour of day
- Looked for trends in network usage.
  - ◆ Usage peak is M-F 10 AM-6PM EST
  - ◆ Usage trough is weekends and early mornings (4 AM-7AM)
- Now collecting this data for longer term trending.
  - ◆ Summer not best time for load, but this did include the EROs




Archive Contribution to Network Usage



# Current State: Archive Contributions to Network Load

- Archive contributes ~30% of network load
- Archive usage trends are different than overall trends.
- Predicting growth rates for non-archive contribution may not be possible.

STDATU Network Contribution  
(8/14/09 - 9/10/09)




# Current State: User Performance

- Users rarely see performance better than 10 Mbps
  - ◆ 1 GB = 00:14:19 @ 10Mbps
  - ◆ Nominal dataset is about 1GB.


**STDATU FTP 2009 External User Performance**

~1.2 M FTP Requests



**STDATU HTTP 2009 User Performance**

~18M HTTP Requests



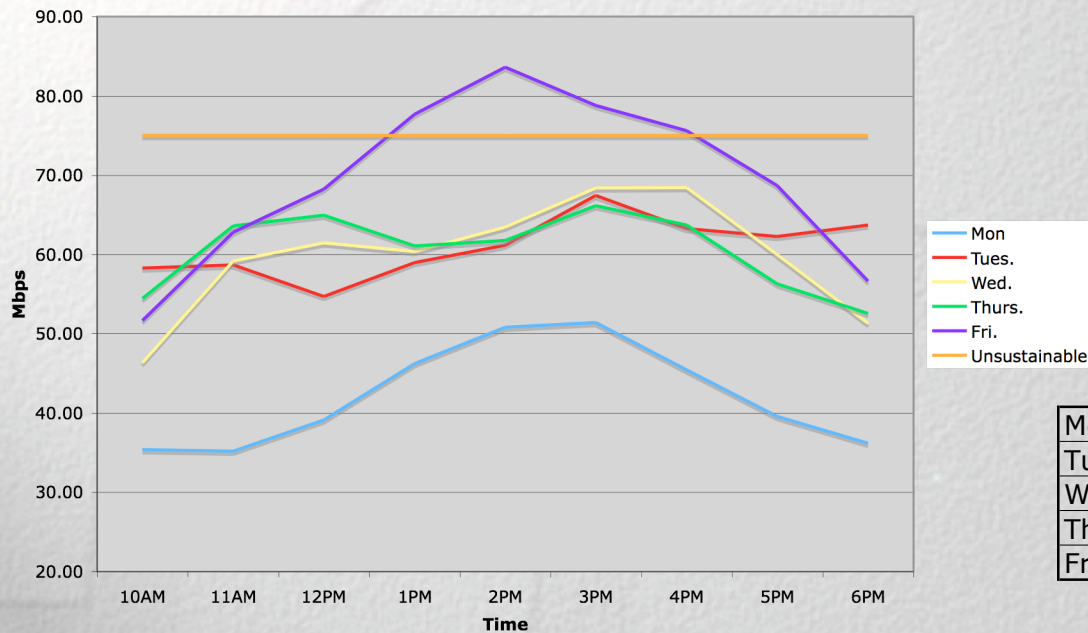


# Current Mitigation Activities

- Archive group is already doing things to reduce network traffic.
  - ◆ HLA transfers being made using external hard drives
    - ◆ Shipping out ~2 TB/Month to CADC and ECF.
      - ~163 hours @ 30 Mbps (but only on weekends)
    - ◆ Shipping in CADC reprocessed WFPC2 data (12 TB).
      - ~40 days @ 30 Mbps (but only on weekends)
  - ◆ GALEX data transferred by hard drives
    - ◆ Recent data releases was ~ 11 TB
    - ◆ Small GO releases are sometimes transferred via network.
  - ◆ KEPLER calibrated products are sent via hard drives
    - ◆ ~750 GB per quarter
- Others
  - ◆ DSS and GSC for Goddard mirrors
    - ◆ Host locally for better performance

# What can we support without upgrade?

Network Loads during Prime  
Time



GB available for transfer

	30%	50%	80%
Mon	37.17	61.95	99.12
Tues.	15.91	26.52	42.43
Wed.	17.10	28.49	45.59
Thurs.	16.39	27.32	43.71
Fri.	6.43	10.72	17.15

- Prime hours can support an increase of about 100 GB to 250 GB per week.
  - ◆ Some days (Fri.) are really bad
  - ◆ Performance to users will likely be worse than current metrics.
- Bandwidth competition will create very unsatisfying experiences for the user.
  - ◆ Ex: Additional request on Friday @ 1PM



# Closing Comments/Thoughts

- Performance will only decrease as volume increases.
  - ◆ GALEX is generating a lot of traffic.
  - ◆ New instruments will begin to go public near term.
  - ◆ Data Storage volume is estimated to increase to 1 PB within the next year.
- Dataset sizes will continue to grow.
- Overall network traffic will continue to grow.
- User performance expectation will continue to grow.
- How do we support mirror site requests? Maintain integrity?
- Archive model (e.g. HLA) moving forward is more interactive with the user.
  - ◆ Using tools to locate data they want as opposed to simply retrieving known datasets.
  - ◆ More tools and services for research
- Archive model will become more active.
  - ◆ Users being notified by the archive when a past dataset has new calibration files.

# Seeking ....

- *Recommendation from this group for STScI to actively pursue increasing or adapting its external network connections to levels that maintain the appropriate support for the user community.*