

Data Transfers in the Grid: Workload Analysis of Globus GridFTP

Nicolas Kourtellis, Lydia Prieto, Gustavo Zarrate, **Adriana Iamnitchi**
University of South Florida

Dan Fraser
Argonne National Laboratory



Globus Software: dev.globus.org

Globus Projects

MPICH G2
OGSA-DAI
Incubation Mgmt

Java Runtime	Delegation	MyProxy	GRAM	Replica Location
C Runtime	CAS	GSI-OpenSSH	GridFTP	MDS4
Python Runtime	C Sec	GridWay	Reliable File Transfer	GT4 Docs

Globus Toolkit

Incubator Projects

			Swift	GEMLCA	gRAVI		MonMan
		GAARDS	MEDICUS	Cog WF	Virt WkSp		NetLogger
GDTE	GridShib	OGRO	UGP	Dyn Acct	Gavia JSC	DDM	Metrics
Introduce	PURSE	HOC-SA	LRMA	WEEP	Gavia MS	SGGC	ServMark

Common Runtime

Security

Execution Mgmt

Data Mgmt

Info Services

Other

Objective 1: Quantify volume of transfers

- What is the transfer size distribution?
- What is the volume of activity for the most active hosts?

Objective 2: Understand how tuning capabilities are used

- What are the buffer sizes used during the transfers?
- What is the average bandwidth?
- What is the utilization of functionalities like streams and stripes?

Objective 3: Quantify user base and predict usage trends

- How does the user base evolve over time?
- What are the geographical characteristics of the GridFTP data transfers?

Outline

- Metrics dataset
- Surprises and ...
- ... zoom in (TeraGrid)
- Lessons and discussions

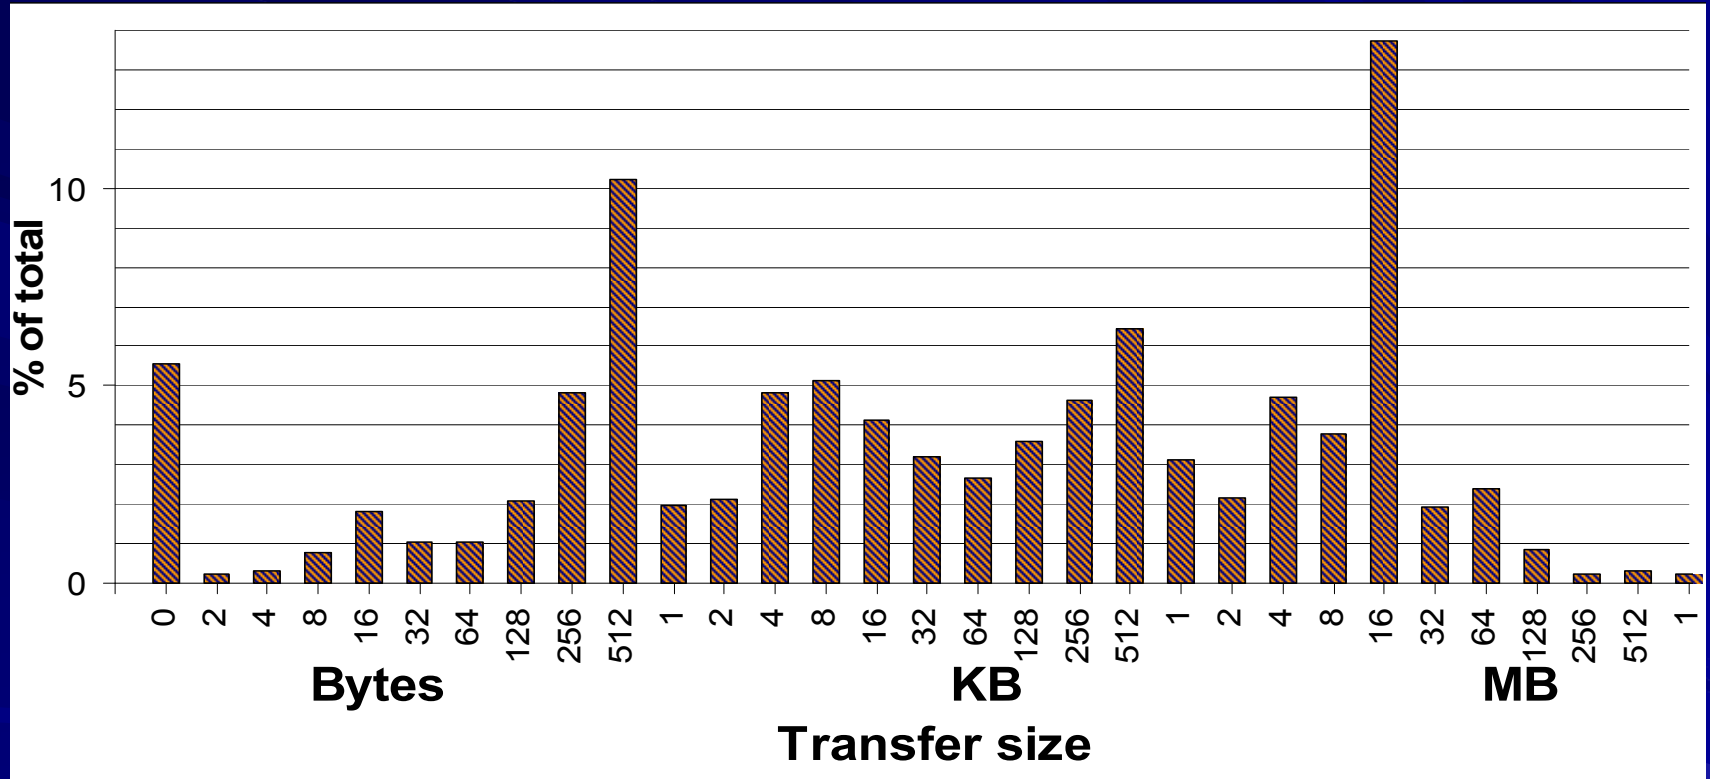
GridFTP Metrics Dataset

Field	Range of Values	Comment
Source hostname/host IP	String/IPnet	Anonymized
Start time of the transfer	Timestamp	Accuracy: ms
End time of the transfer	Timestamp	Accuracy: ms
TCP Buffer Size	Integer (Bytes)	≥ 0
Total Number of Bytes	Integer (Bytes)	≥ 0
Number of Streams	Integer	≥ 1
Number of Stripes	Integer	≥ 1
Store or Retrieve	Integer (0, 1,2)	STOR, RETR, LIST

Metrics Dataset

- Started with ~137.5 million records (Jul'05 - Mar'07)
- Cleaning:
 - transfer size ≤ 0 : ~22.8 million records
 - buffer size < 0 : ~1000 records
 - directory listings: ~3.9 million records
 - invalid hostnames (e.g., /[B@89712e): ~4,600 records
 - ANL-TeraGrid testing: ~11.4 million records
 - duplicate reports: ~16.8 million records
 - self transfers (source=destination): ~5.75 million records
- Clean database: ~77.2 million records (~56.2%)

Surprise #1: Transfer Size Distribution




Objective 1: Quantify volume of transfers

Zoom-in: TeraGrid

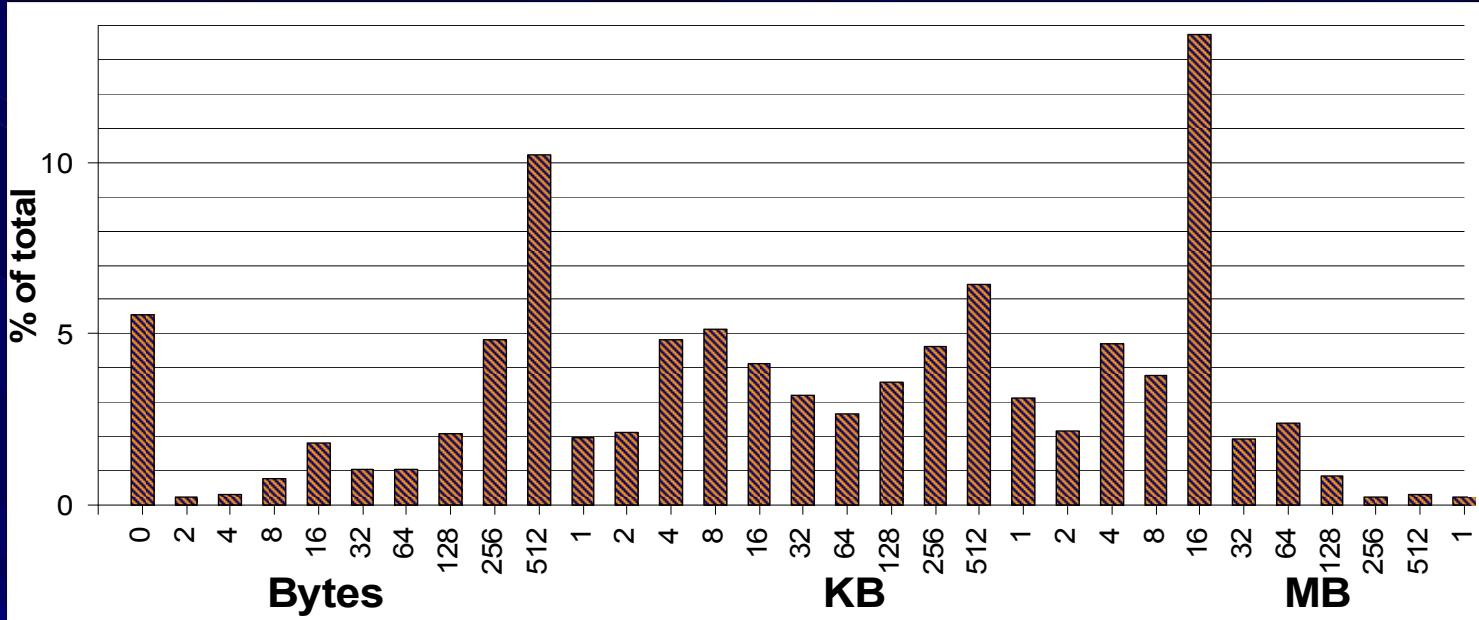
- Are these results representative for production grids?
 - GridFTP testing for deployment and learning
- Identify transfers from TeraGRid and analyze dataset.

Transfer Size Distribution (TG)




Objective 1: Quantify volume of transfers

All



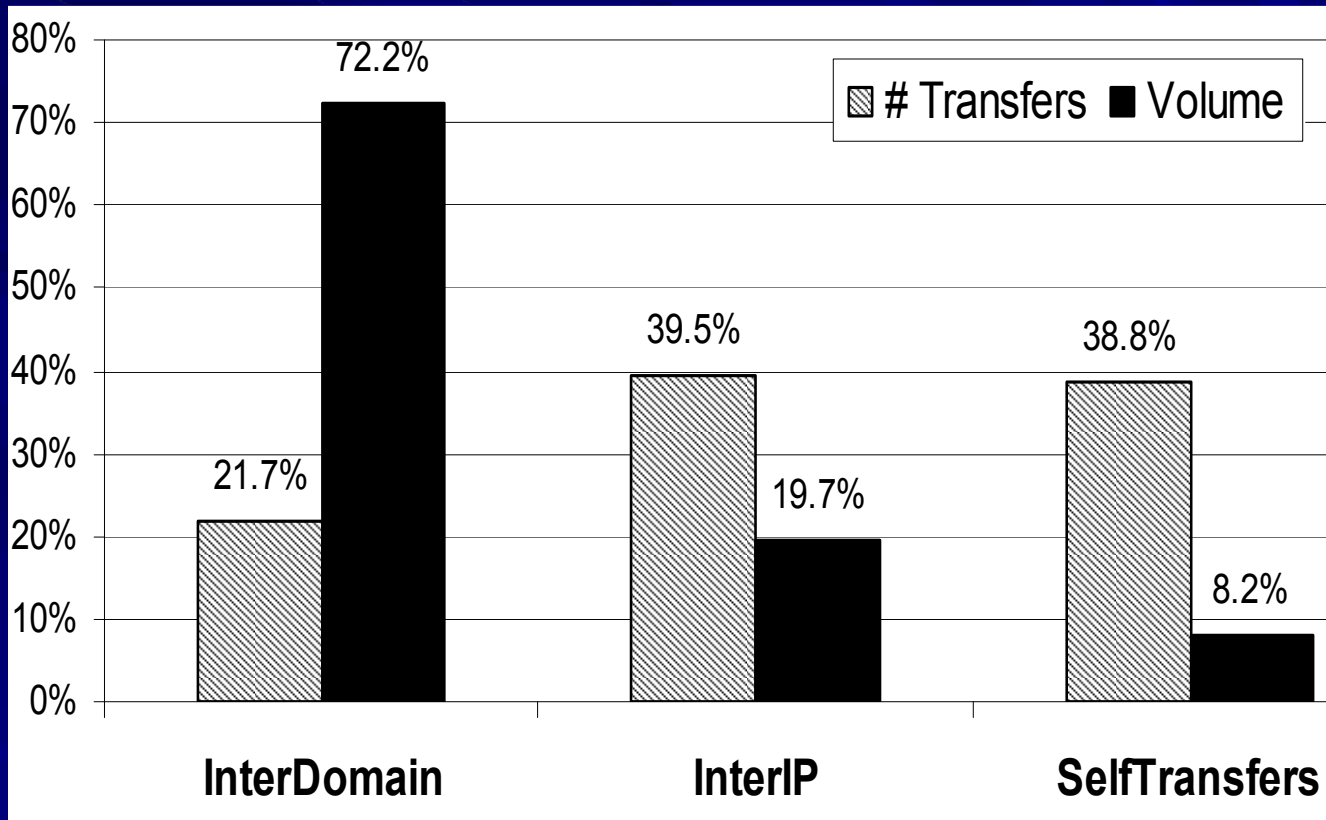
TG



Why So Small Transfers?

- There are still many old versions (i.e., before v3.9.5) of GridFTP in use. These versions do not include trace reporting capabilities.
- Other data transfer protocols and implementations are used
- Users have turned off the reporting capability
- Some of the logs are inevitably lost due to the UDP-based reporting mechanism
- The low transfer volumes could suggest a shift towards data-aware job scheduling (?)

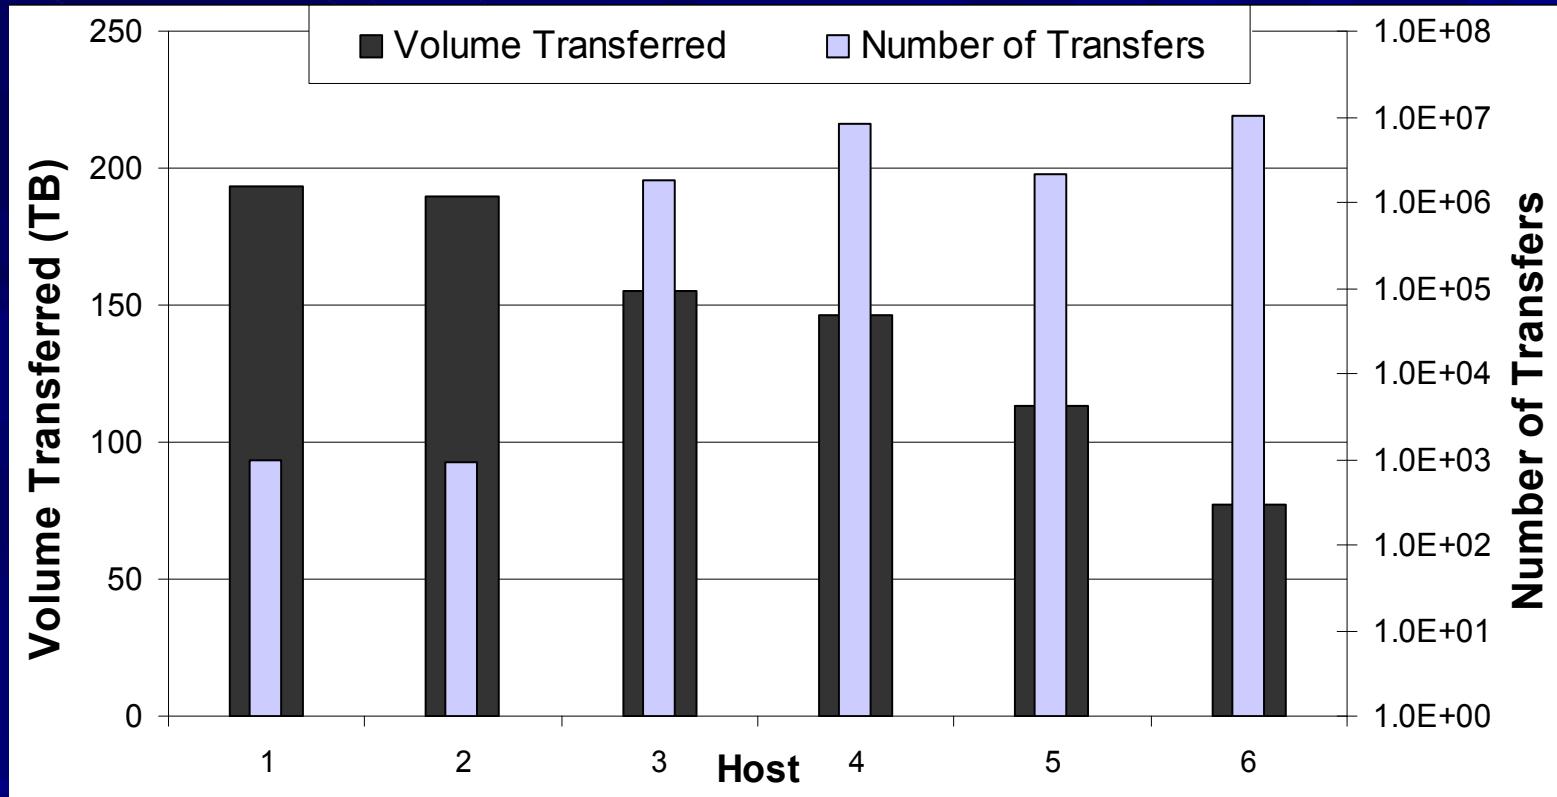
Server to Server Transfers



- High reporting of Self Transfers (more than 1/3)

Objective 1: Quantify volume of transfers

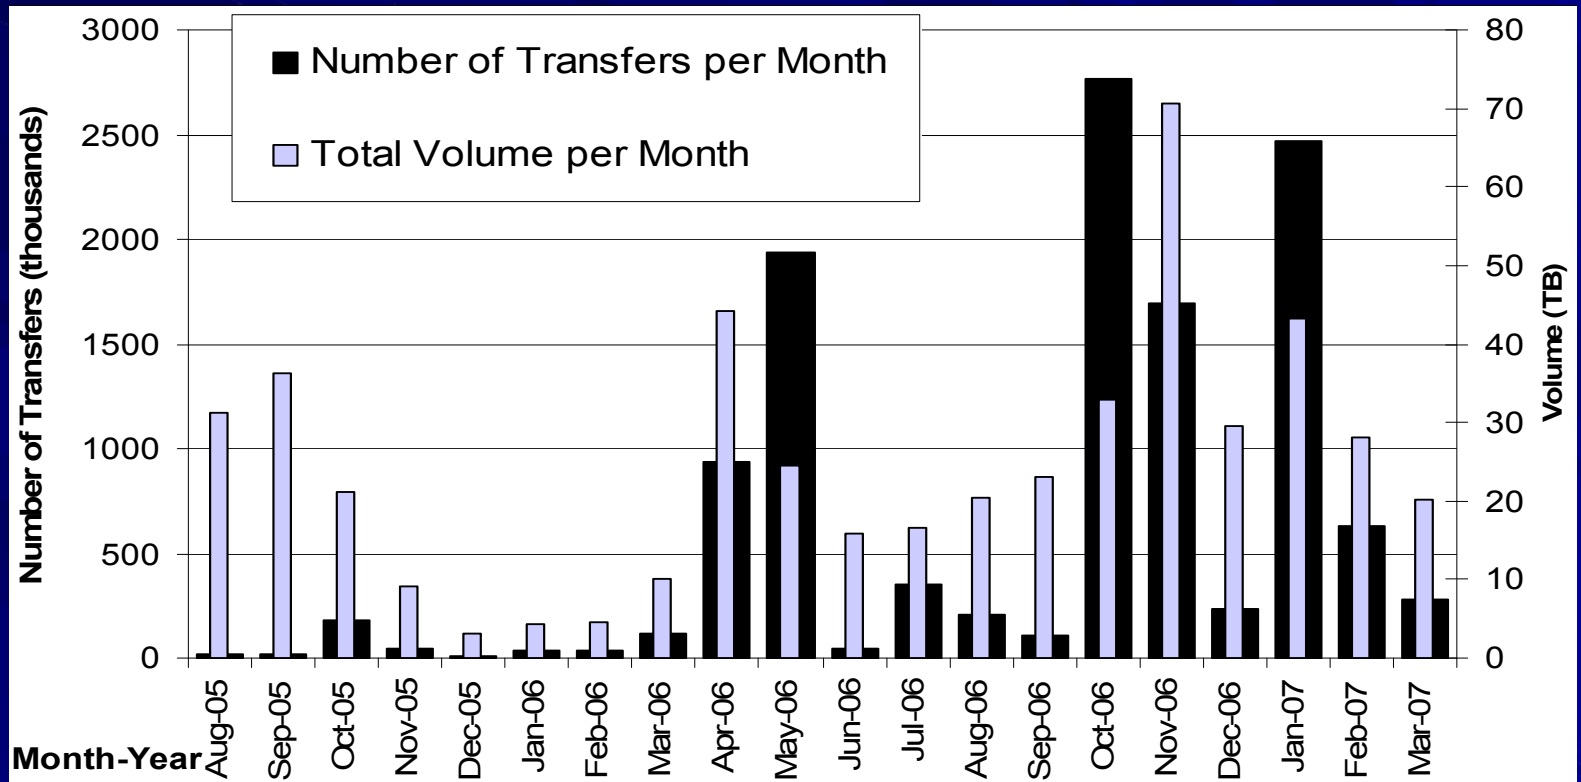
Top 6 Active Hosts (all)



- Top 6 hosts traffic adds up to ~28% of total volume
- Next 48 hosts (IPs) transferred 10s of TB

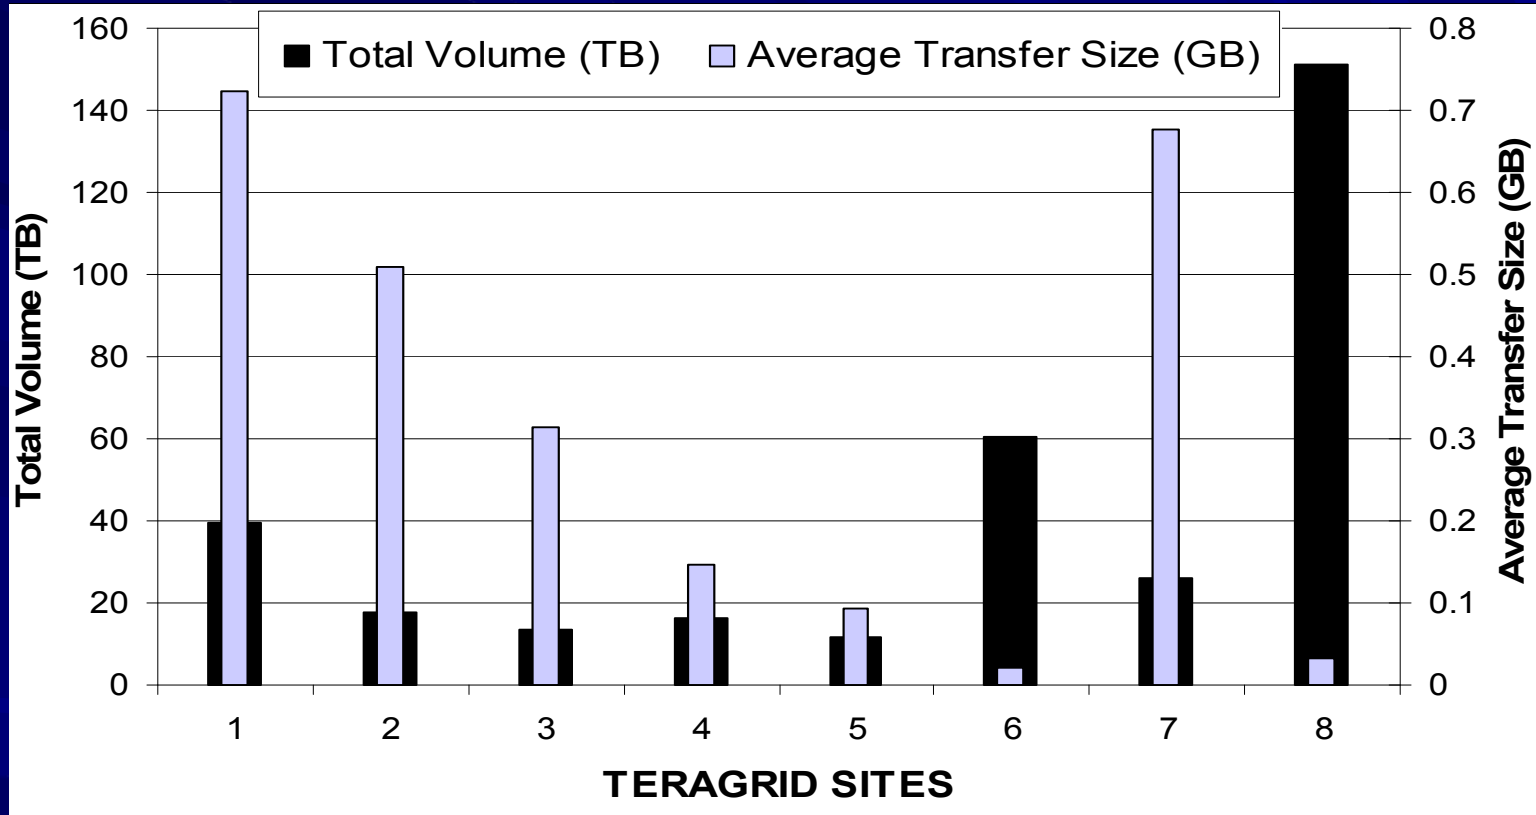
Objective 1: Quantify volume of transfers

Number of Transfers & Volume (TG)



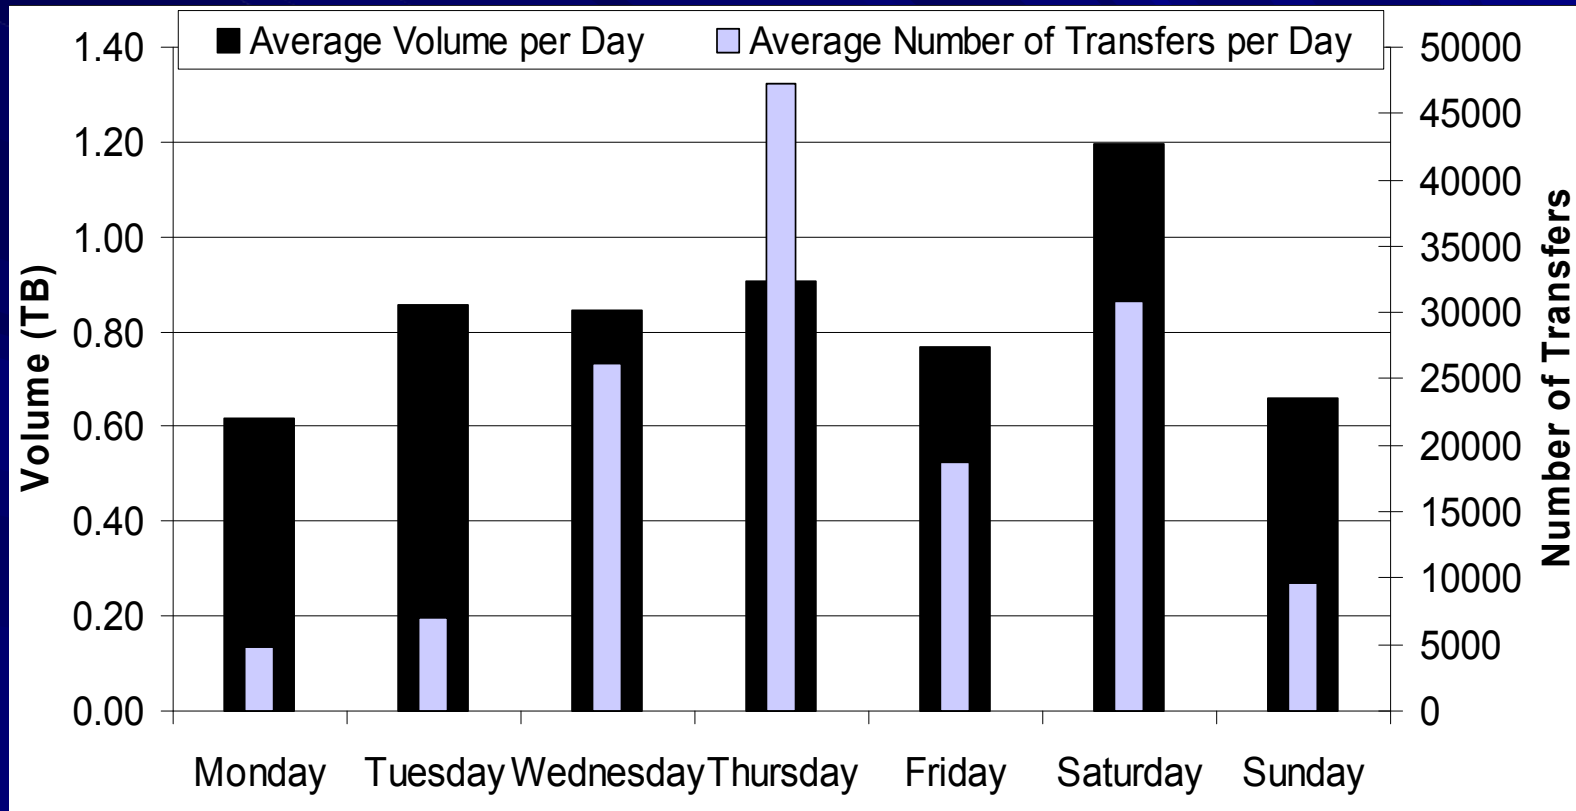
Objective 1: Quantify volume of transfers

Average Transfer Size & Total Volume (TG)



Objective 1: Quantify volume of transfers

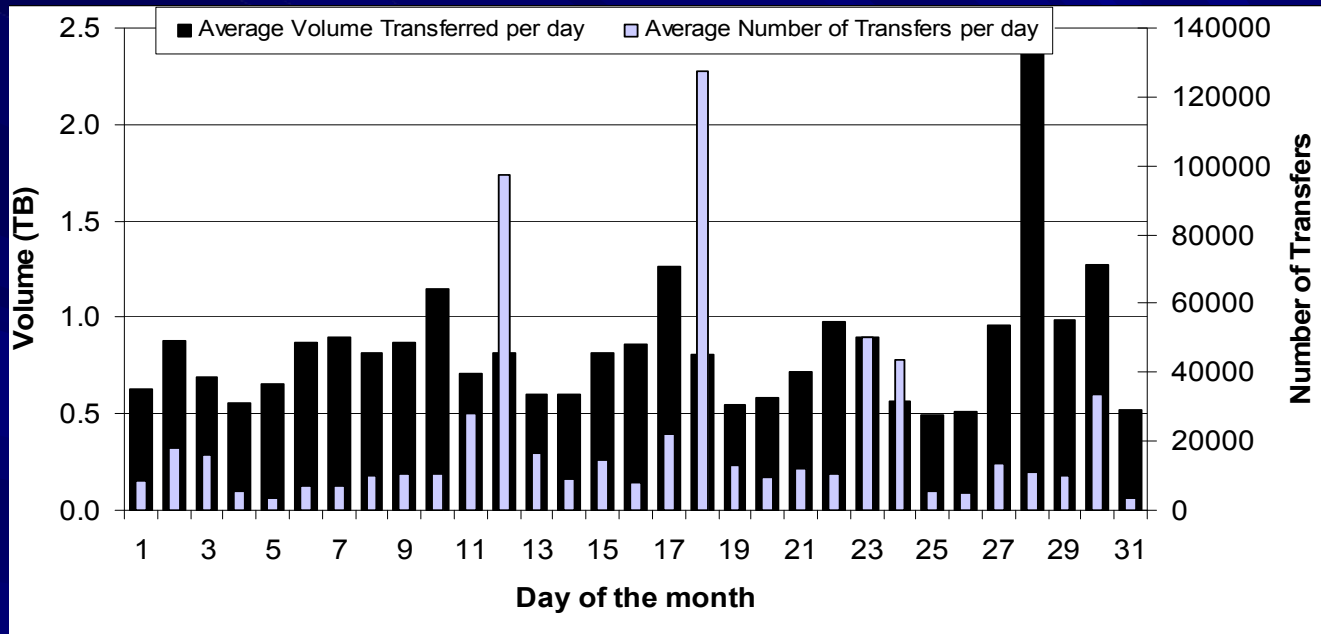
Daily Workload (TG)



- Average volume transferred per day: ~ 0.6TB
- GridFTP doesn't get weekends free!

Objective 1: Quantify volume of transfers

Monthly Workload (TG)



- ~50,000 transfers per day
- ~1TB per day of total volume
- Lowest around 0.5TB per day
- Peaks due to particular days

Objective 1: Quantify volume of transfers

Objective 1: Quantify volume of transfers

- What is the transfer size distribution?
- What is the volume of activity for the most active hosts?

Objective 2: Understand how tuning capabilities are used

- What are the buffer sizes used during the transfers?
- What is the average bandwidth?
- What is the utilization of functionalities like streams and stripes?

Objective 3: Quantify user base and predict usage trends

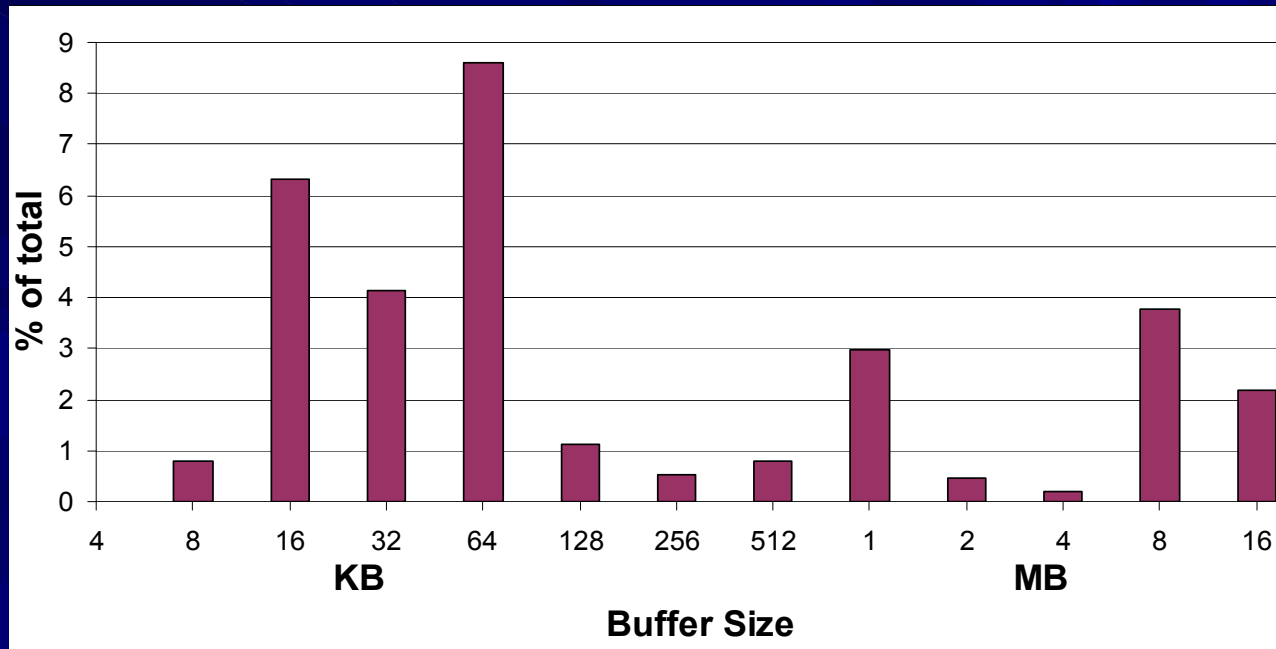
- How does the user base evolve over time?
- What are the geographical characteristics of the GridFTP data transfers?

Surprise #2: Usage of Streams and Stripes

- Unreliably reported (from Globus team).
- Reliable observations:
 - At least ~20% of the transfers used 4 streams (suggested number by ANL's website)
 - At least ~10% of the transfers used a different value, larger than one.
 - Maximum number of streams reported: 1010 (!!)

Objective 2: Understand how tuning capabilities are used

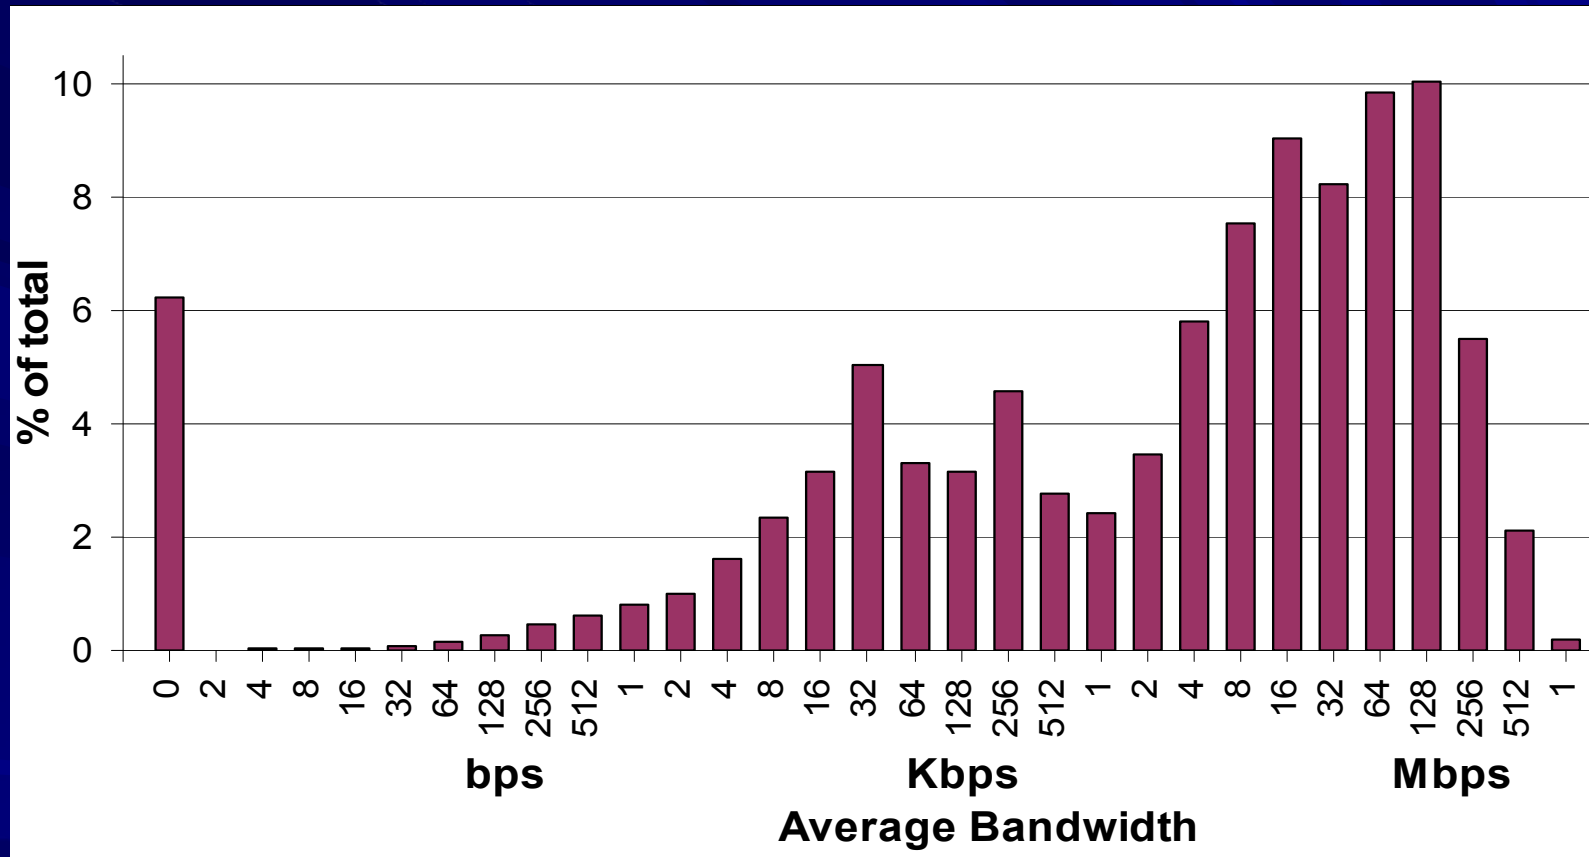
Buffer Size Distribution



- 60% from the original table: OS-controlled (0 bytes)
- Most commonly used: 16–128KB
- Largest buffer size: 1-2GB (92 records)

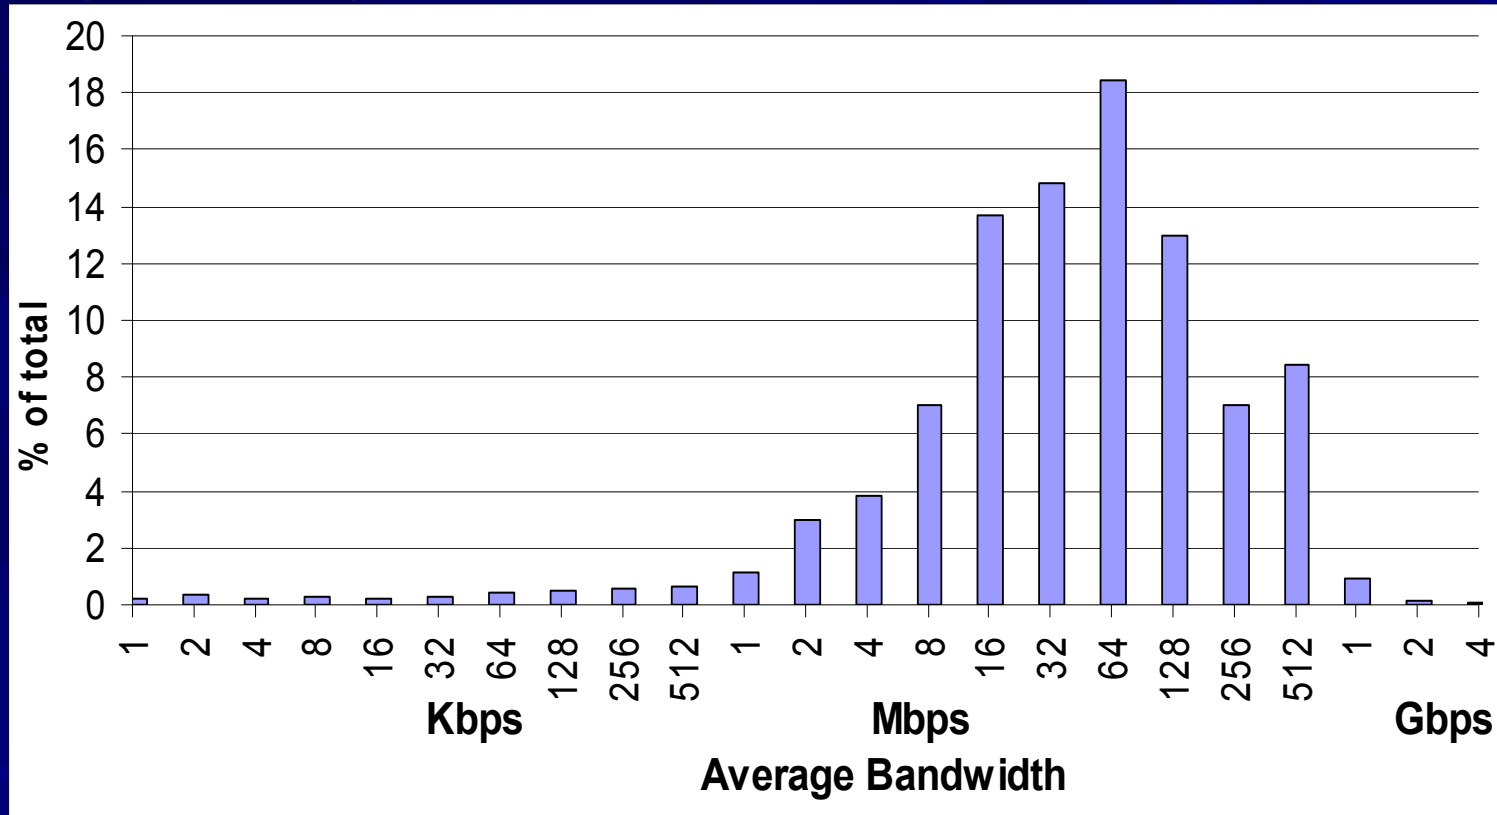
Objective 2: Understand how tuning capabilities are used

Average Bandwidth Distribution



- Peak: 128-256Mbps, ~7.7 million records
- Most common (58%): 4Mbps—1Gbps

Average Bandwidth Distribution (TG)



Compared to the total dataset:

=> The region of 4Mbps to 1Gbps includes more than 85% of the transfers (58% for the whole dataset)

Objective 1: Quantify volume of transfers

- What is the transfer size distribution?
- What is the volume of activity for the most active hosts?

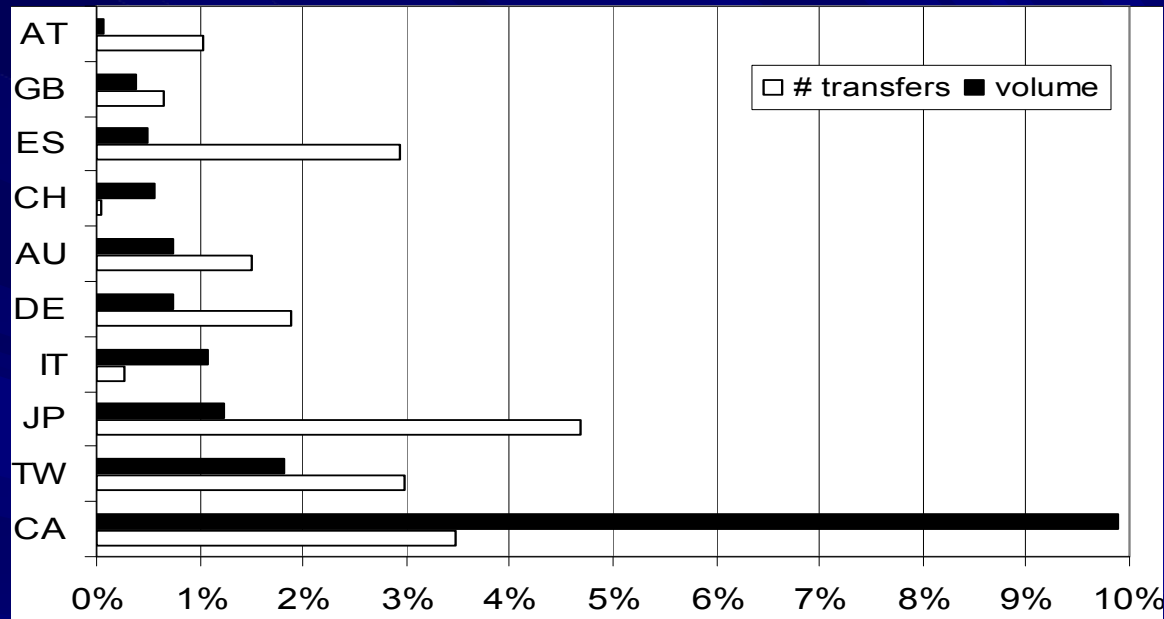
Objective 2: Understand how tuning capabilities are used

- What are the buffer sizes used during the transfers?
- What is the average bandwidth?
- What is the utilization of functionalities like streams and stripes?

Objective 3: Quantify user base and predict usage trends

- How does the user base evolve over time?
- What are the geographical characteristics of the GridFTP data transfers?

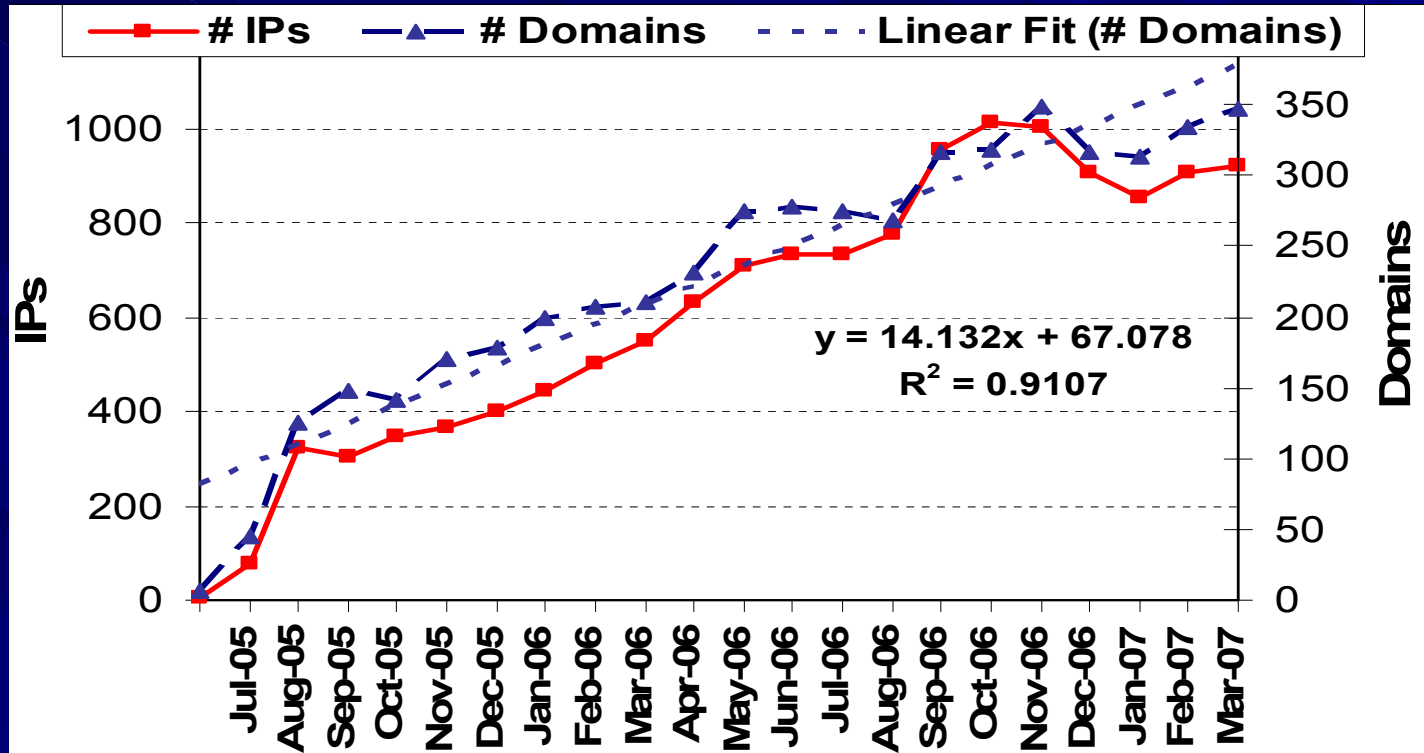
Surprise #3: Geographic Distribution



- USA: 78.4% or ~50.8 million transfers and 82.9% or ~1.7 PB
- Canada+Taiwan+Japan+Spain:~14M transfers and 346TB
- 49 different countries and 446 different cities (178 cities from USA)

Objective 3: Quantify user base and predict usage trends²⁴

User and Domain Evolution (all)



- Continuing increase of user and organization population
- Forecasts: 67 new IPs and 14 new domains per month

Objective 3: Quantify user base and predict usage trends²⁵

Summary of Results

- Many transfers in the range of KBs to 10s MB (peak in 16MB-32MB).
 - relevant for setting up realistic simulations.
 - previous work assume different, larger file sizes.
- Bandwidth measured in previous work is confirmed by our workload analysis.
- Tuning parameters:
 - Users tend not to set the buffer size explicitly (60%), leaving it to the OS
 - The unexpectedly small transfers do not justify tuning GridFTP parameters (stripes and streams)
- The usage of Globus GridFTP is growing over time in terms of IPs (users), domains (organizations), and volume transferred.
- Missing some of the big players

