FLUME
NETWORK OPTIMIZATION

Move and synchronize globally distributed data at unprecedented speeds

Flume™ is an advanced file transfer and synchronization tool that significantly accelerates data transfers across global networks. Traditional TCP/IP based tools work well on local area networks but slow down and frequently abort as distances, congestion and network intermittencies increase. **Flume** works by virtually eliminating the challenges of latency and mitigating the effects of line degradation/congestion that limit traditional IP transfer technologies. By using **Flume**, a company can now move critical data across the internet faster and easier than ever before.

FLUME PROVIDES:

- 5-100x acceleration on challenged connections
- A software only solution that runs on commodity PC hardware
- A solution that runs in virtual, cloud and grid environments
- Acceleration of data without the need to change any of the intermediary network
- Ability to handle and guarantee accuracy of all data types and file formats
- A suite of data efficiency and security features: advanced data synchronization; configurable encryption and compression.

Flume accelerations on world wide networks

Link	Bandwidth	Acceleration
Boston – San Jose	T1	**5x**
China – San Jose	T2	**6x**
India – San Jose	2x T2	**7x**
Israel – India	T3	**38x**
China – Netherlands	OC48	**111x**

Technical summary

Flume is a complete data transfer and synchronization utility that maximizes data throughput by minimizing the effects of latency, intermittency and congestion on worldwide data transfers. The core Flume Engine achieves its increased acceleration and throughput without modifying the data in any way or depending on previous transmissions to serve as foundations for incremental processing. Flume can also leverage optional data efficiency technologies such as block level synchronization and configurable compression to further accelerate file transfer operations. Flume reacts to past network conditions and predicts future network congestion in order to maximize use of available bandwidth while remaining “fair” to competitive traffic as demanded by network managers.

Why are traditional internet communications so poor over long haul networks? **Network throughput using the standard TCP protocol is severely limited by several factors:**

LATENCY

More than half the total overall transmission time for a large file can consist of waiting for data packets to be transmitted to the receiver and acknowledgements of successful receipt to be returned. Due to this latency, increasing transmission bandwidth rapidly reaches a point of diminishing returns (see chart).

NETWORK CONGESTION

Most file transmissions today compete for available network capacity with many other users. TCP adapts very poorly to changing traffic volumes.

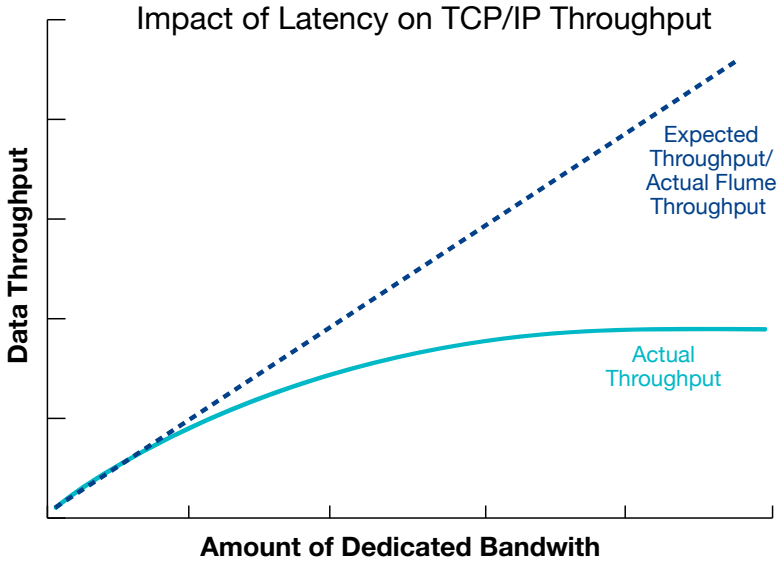

NETWORK QUALITY

Lost or distorted data results in partial retransmission and magnified latency at best and total transmission aborts at worst.

Frequently Asked Questions

What is Flume's architecture? What network investment do I need to make to use Flume?

FLUME is a server-client software solution. Flume does not require any special network hardware. FLUME runs on standard or virtual Linux machines with one machine required on each end of a desired connection.

What speeds can Flume achieve on my network?

FLUME has continually been proven to leverage all available bandwidth - while being respectful to concurrent network traffic. A given network's characteristics (bandwith, congestion, distance) will determine the actual acceleration potential, but Flume has consistently proven to provide acceleration benefits as data is sent outside the local area network.

What application functionality does Flume offer?

FLUME is currently packaged in a configuration similar to the rsync directory synchronization utility with additional feature sets, such as advanced reconnection, data piping and packet level data integrity checking, not included in the rsync application.

CASE STUDY - US Financial Services Company

A large US financial services company with worldwide operations was unable to replicate data on all sites despite having dedicated T1 lines. The company was only able to attain a 35 KB/s rate of throughput on 250 KB lines due to long haul latency. Installing Flume increased throughput to 170- 200 KB/second dependent on concurrent network traffic. The customer added 3 more T1 lines to their network and Flume scaled linearly giving 930-970 KB/s of throughput on the available 1MB/s of bandwidth without interrupting regular IP communications among the sites.

System Requirements:

Hardware Requirements:

Flume 2.2 runs on 32-bit x86 or 64-bit x86_64 Intel or AMD processors with 256MB of RAM .

Operating System Requirements:

32-bit or 64-bit Linux 2.6 kernel: Red Hat Enterprise Linux 4 or 5, CentOS 4 or 5, SUSE Linux Enterprise Server 9 or 10, Ubuntu 7.10

Windows XP / Vista with Linux as above running in a virtual machine

SARATOGA DATA SYSTEMS

Saratoga Data Systems, Inc.

Saratoga is a software company dedicated to solving the problems caused by huge data volumes. The founders all have extensive software development experience applying limited computer resources in the design and manufacturing of data processing applications and the design of state-of-the-art computer integrated circuits.

2530 Berryessa Rd. #424 San Jose, CA 95132 | 408.898.4307 | saratogadata.com