

# Xrootd Monitoring

Atlas Software Week

CERN

November 27 – December 3, 2010

Andrew Hanushevsky, SLAC

---

# Outline

---

- # Introduction to xrootd monitoring
  - What's available
  - How it works
  - What you probably really want

# What is xrootd monitoring?


- # Server-side services that report information
  - Two services configured via the xrootd config file
    - Real-time detailed monitoring
      - xrootd.monitor directive
    - Periodic summary monitoring
      - xrd.report directive
  - Details in “Xrd/Xrootd Configuration Reference”
    - [http://xrootd.org/doc/prod/xrd\\_config.htm](http://xrootd.org/doc/prod/xrd_config.htm)
    - [http://xrootd.org/doc/prod/xrd\\_config.pdf](http://xrootd.org/doc/prod/xrd_config.pdf)

# Why Two Services?

---

- # Real-time and periodic data vastly different
  - Each service designed for unique data requirements
    - Real-time data is fast paced and continuous
      - Setting this up is not for the faint-of-heart
    - Periodic data is rather leisurely but bursty
      - Has few requirements and is relatively easy to setup
- # It's likely you will only use periodic summaries
  - As we shall see as we go on

# Real-Time Monitoring Flow



- **Start Session** `sessionId`, user, pid, client, server, start T, authinfo
- **FRM Staging** stageid, user, pid, client, path, stage T, size, duration, server
- **Open File** fileid, user, pid, client, server, file path, open T
- **File I/O** fileid, I/O length, offset, window T
- **Close File** fileid, bytes read, bytes written, close T
- **App data** user, pid, client, server, application specific data
- **End Session** `sessionId`, duration, end T

# Real-Time Data Handling Is Hard

- # Potentially > 50 MB/Sec monitoring stream
  - Needs fast data collector (i.e. monitoring server)
    - Part of the base package
- # Complex inter-related data
  - Needs sophisticated tools to probe relationships
    - Base package interfaces with MySQL
    - Provides basic web-interface rendering of data
- # This is a lot of work to put up!

# But Can Get Very Good Insights

## Top Performers Table

Table rows: 5 Time Period: Last Hour Site: SLAC Update

| Top active users         |                     |                     |                         |                     |                       |                         |                         |
|--------------------------|---------------------|---------------------|-------------------------|---------------------|-----------------------|-------------------------|-------------------------|
| User Name                | Now                 |                     |                         | Last Hour           |                       |                         |                         |
|                          | Number of Jobs ↑    | Number of Files     | File Size [MB]          | Number of Jobs      | Number of Files       | File Size [MB]          | MB Read                 |
| <a href="#">ayarritu</a> | <a href="#">615</a> | <a href="#">139</a> | <a href="#">65,987</a>  | <a href="#">430</a> | <a href="#">146</a>   | <a href="#">65,802</a>  | <a href="#">41,360</a>  |
| <a href="#">iregens</a>  | <a href="#">360</a> | <a href="#">405</a> | <a href="#">371,874</a> | <a href="#">64</a>  | <a href="#">317</a>   | <a href="#">303,252</a> | <a href="#">143,852</a> |
| <a href="#">cschill</a>  | <a href="#">281</a> | <a href="#">32</a>  | <a href="#">27,133</a>  | <a href="#">79</a>  | <a href="#">30</a>    | <a href="#">25,301</a>  | <a href="#">4,892</a>   |
| <a href="#">feltresi</a> | <a href="#">149</a> | <a href="#">106</a> | <a href="#">167,528</a> | <a href="#">70</a>  | <a href="#">143</a>   | <a href="#">218,873</a> | <a href="#">74,552</a>  |
| <a href="#">torsten</a>  | <a href="#">72</a>  | <a href="#">99</a>  | <a href="#">83,673</a>  | <a href="#">184</a> | <a href="#">1,532</a> | <a href="#">630,092</a> | <a href="#">235,327</a> |

| Hottest dataTypes       |                     |                       |                           |                    |                     |                       |                         |                   |                         |
|-------------------------|---------------------|-----------------------|---------------------------|--------------------|---------------------|-----------------------|-------------------------|-------------------|-------------------------|
| dataType Name           | Now                 |                       |                           |                    | Last Hour           |                       |                         |                   |                         |
|                         | Number of Jobs ↑    | Number of Files       | File Size [MB]            | Number of Users    | Number of Jobs      | Number of Files       | File Size [MB]          | Number of Users   | MB Read                 |
| <a href="#">SPskims</a> | <a href="#">998</a> | <a href="#">739</a>   | <a href="#">632,651</a>   | <a href="#">11</a> | <a href="#">663</a> | <a href="#">340</a>   | <a href="#">304,938</a> | <a href="#">6</a> | <a href="#">120,728</a> |
| <a href="#">SP</a>      | <a href="#">652</a> | <a href="#">1,839</a> | <a href="#">1,961,610</a> | <a href="#">12</a> | <a href="#">981</a> | <a href="#">506</a>   | <a href="#">474,819</a> | <a href="#">7</a> | <a href="#">159,512</a> |
| <a href="#">PRskims</a> | <a href="#">93</a>  | <a href="#">650</a>   | <a href="#">811,152</a>   | <a href="#">7</a>  | <a href="#">204</a> | <a href="#">83</a>    | <a href="#">107,807</a> | <a href="#">2</a> | <a href="#">62,265</a>  |
| <a href="#">PR</a>      | <a href="#">66</a>  | <a href="#">600</a>   | <a href="#">453,640</a>   | <a href="#">6</a>  | <a href="#">265</a> | <a href="#">1,454</a> | <a href="#">625,498</a> | <a href="#">3</a> | <a href="#">174,754</a> |
| <a href="#">cfg</a>     | <a href="#">0</a>   | <a href="#">0</a>     | <a href="#">0</a>         | <a href="#">0</a>  | <a href="#">8</a>   | <a href="#">1</a>     | <a href="#">7</a>       | <a href="#">1</a> | <a href="#">10</a>      |

| Hottest skims                     |                     |                     |                         |                   |                     |                     |                         |                   |                        |
|-----------------------------------|---------------------|---------------------|-------------------------|-------------------|---------------------|---------------------|-------------------------|-------------------|------------------------|
| skim Name                         | Now                 |                     |                         |                   | Last Hour           |                     |                         |                   |                        |
|                                   | Number of Jobs ↑    | Number of Files     | File Size [MB]          | Number of Users   | Number of Jobs      | Number of Files     | File Size [MB]          | Number of Users   | MB Read                |
| <a href="#">BtoRhoGamma</a>       | <a href="#">591</a> | <a href="#">139</a> | <a href="#">65,987</a>  | <a href="#">1</a> | <a href="#">458</a> | <a href="#">146</a> | <a href="#">65,802</a>  | <a href="#">1</a> | <a href="#">41,360</a> |
| <a href="#">DstToD0PIToVGamma</a> | <a href="#">262</a> | <a href="#">86</a>  | <a href="#">33,138</a>  | <a href="#">1</a> | <a href="#">70</a>  | <a href="#">41</a>  | <a href="#">16,171</a>  | <a href="#">1</a> | <a href="#">4,668</a>  |
| <a href="#">BToDinu</a>           | <a href="#">115</a> | <a href="#">118</a> | <a href="#">186,026</a> | <a href="#">2</a> | <a href="#">125</a> | <a href="#">145</a> | <a href="#">222,200</a> | <a href="#">2</a> | <a href="#">74,568</a> |
| <a href="#">AllEvents</a>         | <a href="#">76</a>  | <a href="#">394</a> | <a href="#">508,309</a> | <a href="#">3</a> | <a href="#">210</a> | <a href="#">84</a>  | <a href="#">108,365</a> | <a href="#">3</a> | <a href="#">62,268</a> |
| <a href="#">Tau11</a>             | <a href="#">4</a>   | <a href="#">95</a>  | <a href="#">130,103</a> | <a href="#">1</a> | <a href="#">3</a>   | <a href="#">6</a>   | <a href="#">149</a>     | <a href="#">0</a> | <a href="#">127</a>    |

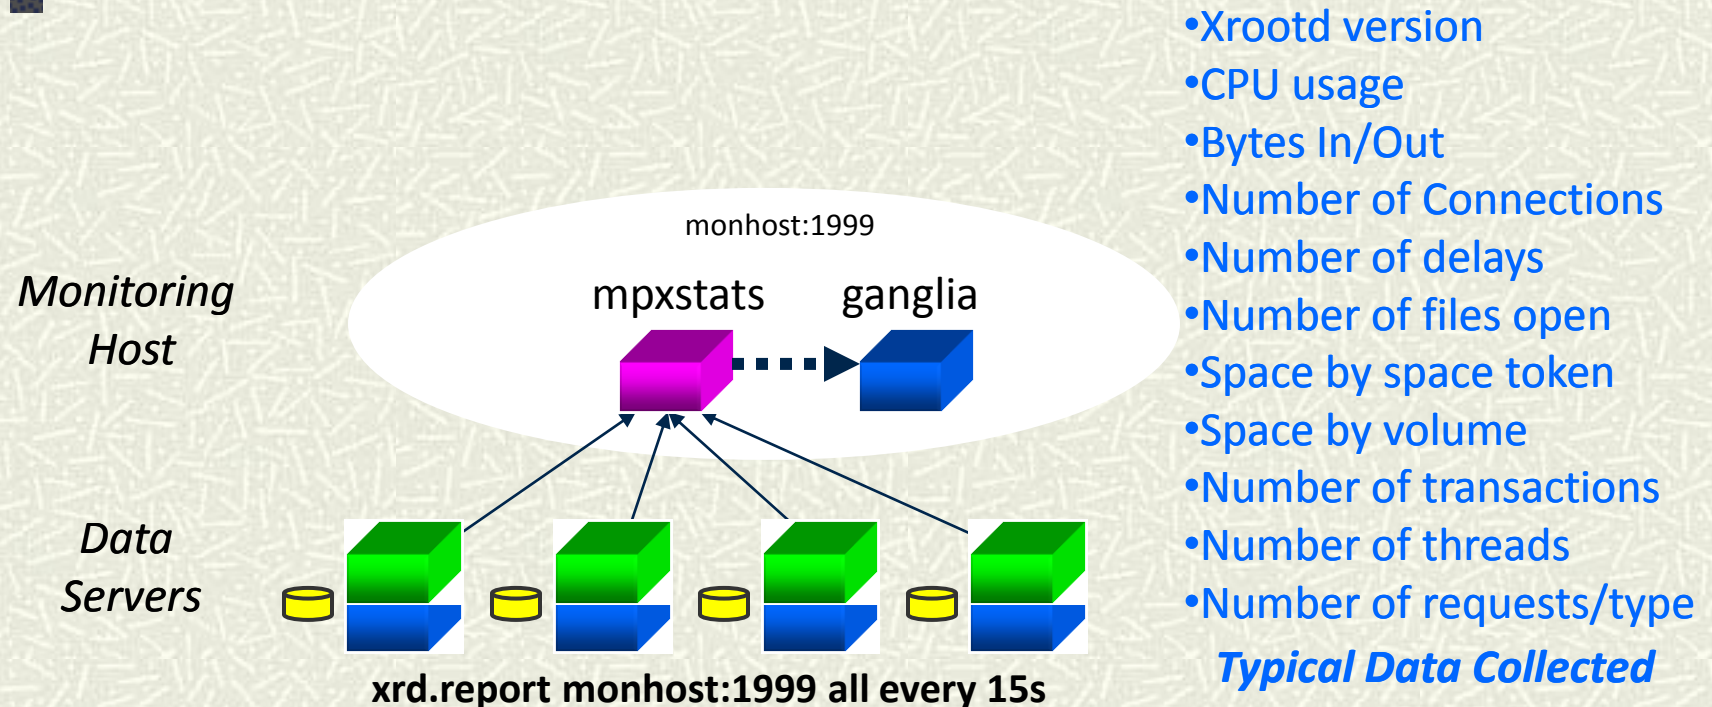
| Hottest files   |                       |                   |                    |                       |
|---|-----------------------|-------------------|--------------------|-----------------------|
| File Path   | File Size [MB]        | Now               | Last Hour          |                       |
|   |                       | Number of Jobs ↑  | Number of Jobs     | MB Read               |
| <a href="#">/store/PRskims/R18/18.6.3d/AllEvents/00/AllEvents_20006.04HB.root</a> | <a href="#">1,690</a> | <a href="#">2</a> | <a href="#">15</a> | <a href="#">1,630</a> |
| <a href="#">/store/PRskims/R18/18.6.3e/AllEvents/05/AllEvents_20502.04HB.root</a> | <a href="#">1,688</a> | <a href="#">1</a> | <a href="#">17</a> | <a href="#">1,636</a> |
| <a href="#">/store/PRskims/R18/18.6.3e/AllEvents/05/AllEvents_20502.01.root</a>   | <a href="#">1,689</a> | <a href="#">1</a> | <a href="#">17</a> | <a href="#">1,635</a> |
| <a href="#">/store/PRskims/R18/18.6.3e/AllEvents/05/AllEvents_20500.03HB.root</a> | <a href="#">1,688</a> | <a href="#">1</a> | <a href="#">19</a> | <a href="#">1,641</a> |
| <a href="#">/store/PRskims/R18/18.6.3e/AllEvents/05/AllEvents_20500.01.root</a>   | <a href="#">1,689</a> | <a href="#">1</a> | <a href="#">19</a> | <a href="#">1,640</a> |

# Summary Monitoring Is Easier

- # Summary data periodically reported
  - Very large amount of data available
    - [http://xrootd.org/doc/prod/xrd\\_monitoring.htm](http://xrootd.org/doc/prod/xrd_monitoring.htm)
  - You pick which is to be reported by category
    - Use the `xrd.report` directive
  - Centrally collect it via provided **mpxstats** tool
    - Merges and converts xml streams to keyword/value pairs
  - Feed data into your favorite monitoring system
    - Active MQ, Ganglia, GRIS, Nagios, MonALISA, etc


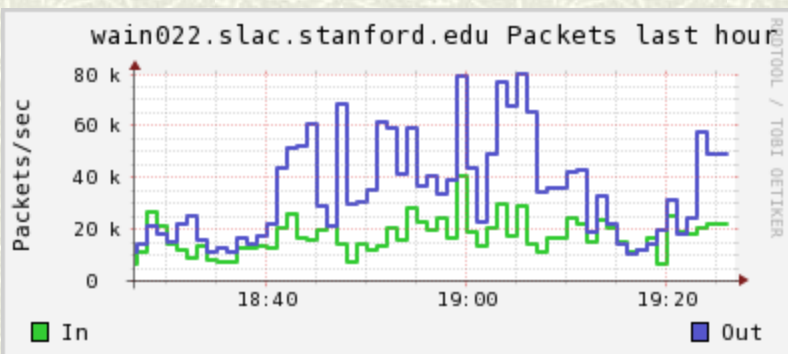
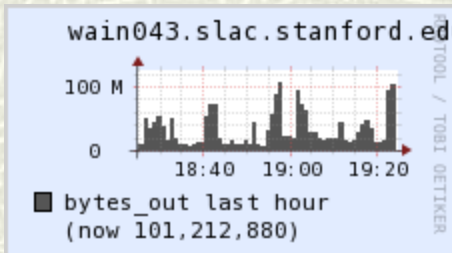

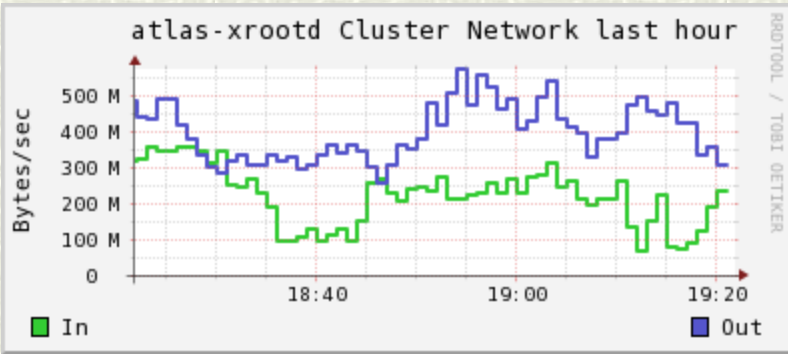


# Summary Monitoring Data Flow



This is a centralized data flow; you can also do a distributed flow!  
Here mpxstats runs on each data server and you route data to localhost.

# Easy To Render Basic Metrics



# Summary

---

- # xrootd provides a wealth of monitoring data
  - From super detailed to basic summaries
- # Your needs will determine what you collect
  - We suggest sticking with periodic summary data
- # However, you must have a monitoring system
  - Ganglia, Gris, Nagios, MonaLisa, or other

# Acknowledgements

## # Software Contributors

- Alice: Derek Feichtinger
- ATLAS: Charles Waldman, Wei Yang
- CERN: Fabrizio Furano, Lukasz Janyst, Andreas Peters
- CMS: Brian Bockelman
- Fermi/GLAST: Tony Johnson (Java)
- FZK: Artem Trunov
- LBNL: Alex Sim, Junmin Gu, Vijaya Natarajan (BeStMan team)
- LSST: Daniel Wang
- Root: Gerri Ganis, Beterand Bellenet, Fons Rademakers
- SLAC: Tofigh Azemmoon, Andrew Hanushevsky, Wilko Kroeger

## # Operational Collaborators

- BNL, CERN, FZK, IN2P3, OSG, SLAC, UC, UTA, UVIC, UWisc

## # Partial Funding

- US Department of Energy
  - Contract DE-AC02-76SF00515 with Stanford University