


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.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, file path, stage T, 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
ayarritu	615	139	65,987	430	146	65,802	41,360
iregens	360	405	371,874	64	317	303,252	143,852
cschill	281	32	27,133	79	30	25,301	4,892
feltresi	149	106	167,528	70	143	218,873	74,552
torsten	72	99	83,673	184	1,532	630,092	235,327

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
SPskims	998	739	632,651	11	663	340	304,938	6	120,728
SP	652	1,839	1,961,610	12	981	506	474,819	7	159,512
PRskims	93	650	811,152	7	204	83	107,807	2	62,265
PR	66	600	453,640	6	265	1,454	625,498	3	174,754
cfg	0	0	0	0	8	1	7	1	10

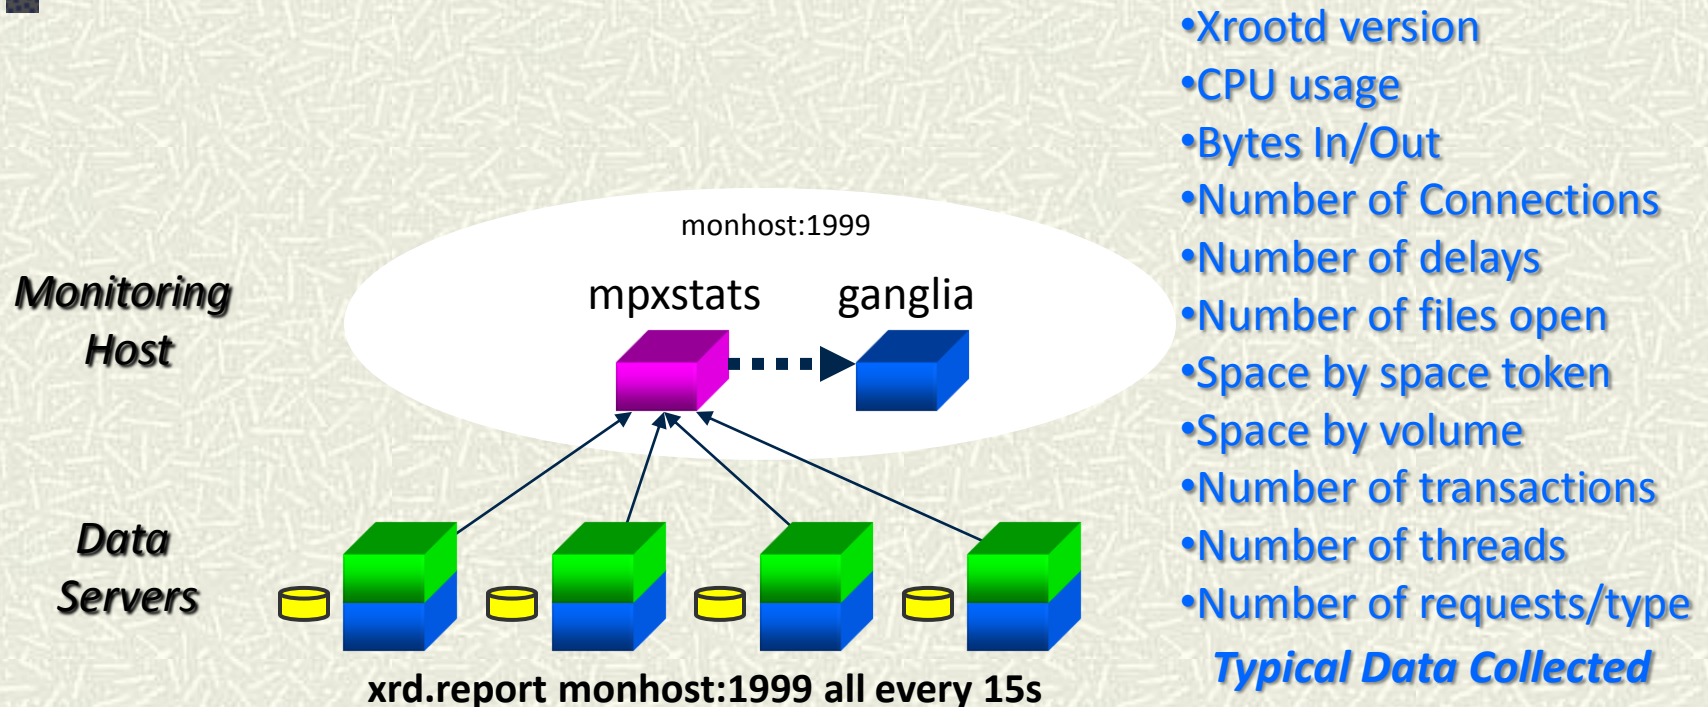
Hottest skims									
skim Name	Now				Last Hour				
	Number of Jobs ↑	Number of Files	File Size [MB]	Number of Users	Number of Jobs	Number of Files	File Size [MB]	Number of Users	MB Read
BtoRhoGamma	591	139	65,987	1	458	146	65,802	1	41,360
DstToD0PIToVGamma	262	86	33,138	1	70	41	16,171	1	4,668
BToDinu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

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

Summary Monitoring Is Easier






- # Summary data periodically reported
 - Very large amount of data available
 - 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
 - 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