



PD2P

The DA Perspective


Kaushik De

Univ. of Texas at Arlington

S&C Week, CERN


Nov 30, 2010

Introduction



- PD2P was introduced ~5 months ago
- PD2P adiabatically and dramatically changed data distribution model for analysis users
 - Designed to have no impact on user experience
 - Changes were rolled in gradually behind the scenes
 - So far, no complaints (or reactions from users) – maybe because they do not know about the changes we made!
- In this talk, I will present
 - Brief overview (for more details see my talk yesterday: <http://indico.cern.ch/getFile.py/access?contribId=14&sessionId=21&resId=1&materialId=slides&confId=76896>)
 - Analysis of usage patterns

Huge Rise in Analysis Activity

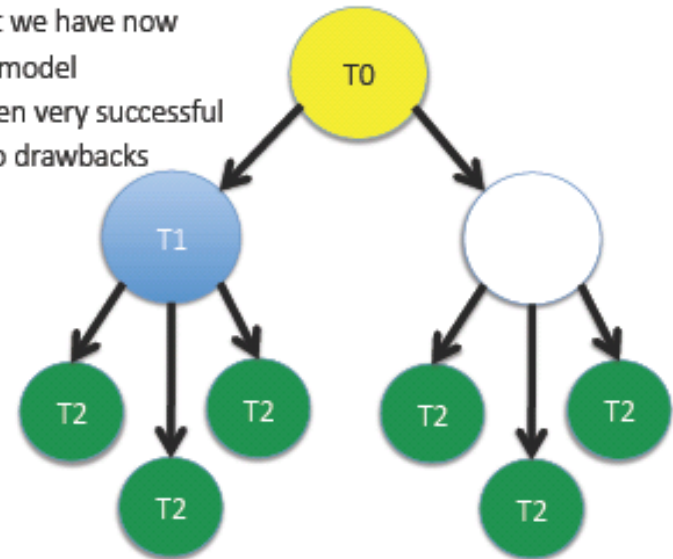


Data Distribution is Very Important

- Most user analysis jobs run at Tier 2 sites
 - Jobs are sent to data
 - We rely on pushing data out to Tier 2 sites promptly
 - Difficult since many data formats and many sites
 - We adjusted frequently the number of copies and data types in April & May
 - But Tier 2 sites were filling up too rapidly, and user pattern was unpredictable
 - Most datasets copied to Tier 2's were never used

Data placement model The "Monarch Model"

- This is what we have now
- It is a push model
- And has been very successful
- But has also drawbacks



We Changed Data Distribution Model

- Reduce pushed data copies to Tier 2's


- Only send small fraction of AOD's automatically
- Pull all other data types, when needed by users
- Note: for production we have always pulled data as needed

- But users were insulated from this change

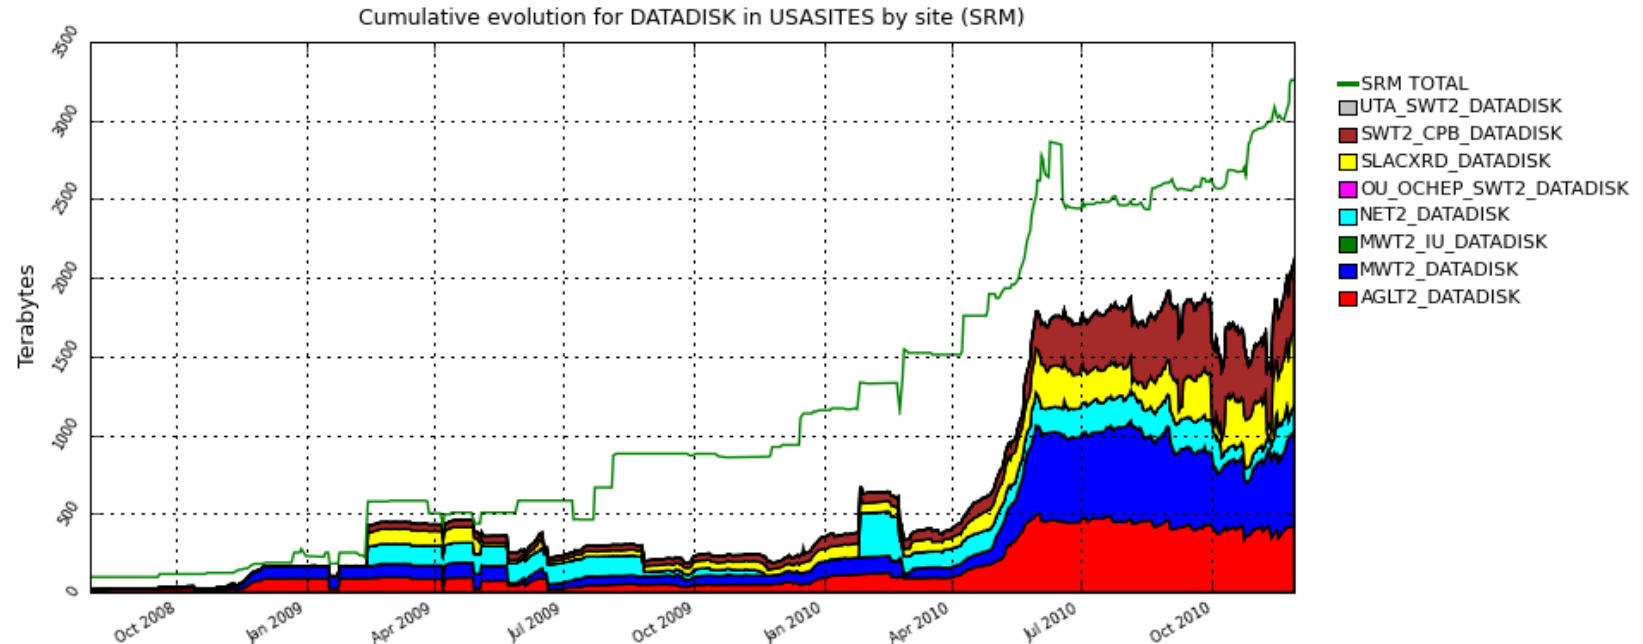
- Did not affect the many critical ongoing analyses
- No delays in running jobs
- No change in user workflow

Data pull model I

- This is Kaushik's PD2P
- Runs now in the US cloud
- Not for RAW and HITS yet
- Interesting results shown



Data Flow to Tier 2's




- Example above is from US Tier 2 sites
 - Exponential rise in April and May, after LHC start
 - We changed data distribution model end of June – PD2P
 - Much slower rise since July, even as luminosity grows rapidly


PD2P Subscriptions



of Datasets Subscribed / 5 days




of Datasets Reused / 5 days




Reuse of Files



of Datasets Subscribed / 5 days




of files from datasets reused / 5 days




Patterns of Data Usage – Part I

- Interesting patterns are emerging by type of data
 - LHC data reused more often than MC data – not unexpected

Reuse of Datasets




Reuse of Datasets




Patterns of Data Usage – Part 2

- Interesting patterns also by format of data
- During past ~5 months:
 - All types of data showing up: ESD, NTUP, AOD, DED most popular
 - But highest reuse (counting files): ESD, NTUP

of Datasets Subscribed by Type




Reuse of Datasets by Type




Trends in Data Reuse

- PD2P pull model does not need a priori assumption about which data types are needed for user analysis
- It automatically moves data based on user workflow
- We observe now a shift back to ESD!

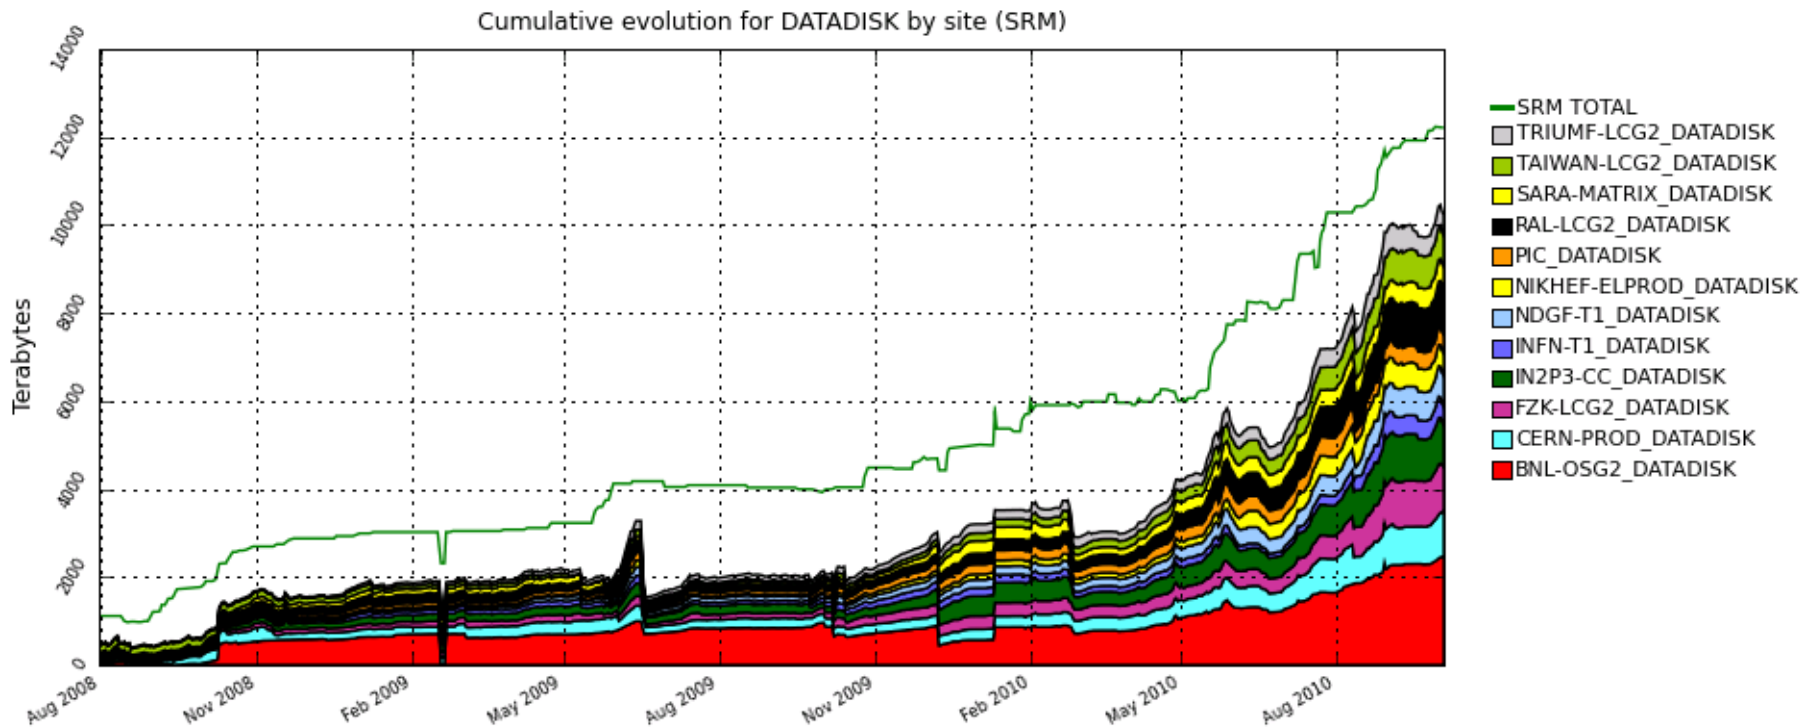


Plans for the Future



- Continue tuning PD2P, as needed
- Develop PD2P for Tier 1's
 - Currently, all data is pushed to Tier 1's through central subscriptions
 - Storage is filling up rapidly – disk crises last month
 - We need to reduce amount of data automatically subscribed
 - Start pulling data as needed for user analysis within cloud
 - We are working out the details of the algorithm
 - More news soon
- Please provide user feedback if something can be improved

Rise in Tier 1 Storage Use



How to Reduce Data Stored at Tier 1's

- Largest amount of Tier 1 data is primary
 - Must reduce automatic subscriptions
- Second largest Tier 1 data is secondary
 - We suspect much of this data is never used
 - Go to pull model for secondary data (and also part of data which is currently labeled primary) using PD2P

