





### Problem

- Modern flle systems can store petabytes of data and billions of files
- Users waste time finding and organizing files
- Administrator decisions are more complicated
- Need a way to quickly find and understand the data being stored

## Searching file metadata

- Inode fields and extended attributes contain useful information
- Can address many user questions
- "Where are the document files I modified most recently?"
- "Which files in my home directory should be deleted?"
- Can address many administrator questions
- "Which system configuration files were changed last week?"
- "Which user home directories have grown most in the past 6 months?" Existing file search tools lack required scalability
- Focus on content-based search, rely on DBMS for metadata search - Content search cannot address many important questions
- DBMS are ill-suited for metadata characteristics

## **Query characteristics**

- File system namespace locality
- Files that satisfy a query are often clustered in the namespace - Due to clustering of metadata attribute values
- Locality of reference
- Like files, only a small set of namespace locations are popular
- Most namespace locations are infrequently searched
- Selectivity
- Queries should try to refine results to a small set of files
- Too many results are not focused and thus not very useful

# **Metadata Characteristics**

We analyzed metadata traces collected from file servers within NetApp 
 Data set
 Description
 # of files
 Server Capacity

 Web
 web & wiki server
 15 million
 1.28 TB

 Eng
 build space
 60 million
 30 GB

 Home
 home directories
 300 million
 76.78 TB

Spatial locality of attribute values:

- Attributes values tend to be clustered in a few sub-trees • Locality ratio: The percentage of directories that contain files with a
- particular attribute value compared to all directories
- Ratios for the ext, size, owner and ctime for the 32 most frequent attribute values

		ext	size	uid	ctime
	Web	0.000162% - 0.120%	0.0579% - 0.177%	0.000194% - 0.0558%	0.000291% - 0.0105%
	Eng	0.00101% - 0.264%	0.00194% - 0.462%	0.000578% - 0.137%	0.000453% - 0.0103%
	Home	0.000201% - 0.491%	0.0259% - 0.923%	0.000417% - 0.623%	0.000370% - 0.128%

 All ratios are below 1% meaning these attributes occur in fewer than 1% of directories

Skewed distribution of attribute values:

- Attribute value distributions tend to be highly skewed
- Data distribution: The highest ranked (most popular) attribute values are plotted against their file count percentage (log-log scale)





- 80% of files have one of 20 popular extensions, while the other 20% account for the other 4000 extensions, following a power law distribution • The cartesian product of top 20 ext and size values is very small
- Attribute combinations are more evenly distributed

# Conclusions

- Growing amounts of data are making data management more difficult Metadata search can address many data management challenges
- Spyglass exploits metadata characteristics for scalable metadata search Achieves up to three orders of magnitude performance improvement compared to DBMS solutions

