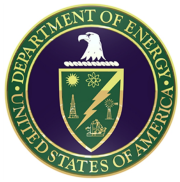


SciDAC
Scientific Discovery through Advanced Computing

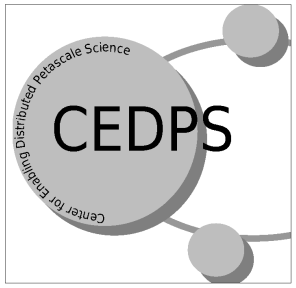


# Distributed Systems Troubleshooting and Performance Analysis

Brian L. Tierney, Dan Gunter

BLTierney@lbl.gov, DKGunter@.lbl.gov

**Distributed Systems Department**
**Lawrence Berkeley National Laboratory**
**http://dsd.lbl.gov**
**Center for Enabling Distributed Petascale Science**
**http://cedps.net**

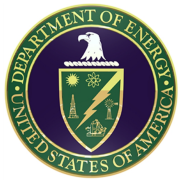


SciDAC
Scientific Discovery through Advanced Computing

# The Problem

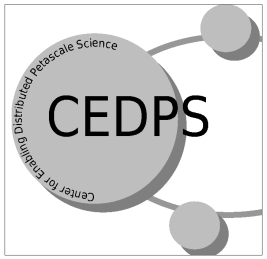


- Assume your distributed compute job normally takes 30 minutes to complete. But…
  - 3 hours have passed and the job has not yet completed.
- What, if anything, is wrong?
  - Is the job still running or did one of the software components crash?
  - Is the network particularly congested?
  - Is the CPU particularly loaded?
  - Is there a disk problem?
  - Was a software library containing a bug installed somewhere?


SciDAC
Scientific Discovery through Advanced Computing

# Distribution of Job completion Time

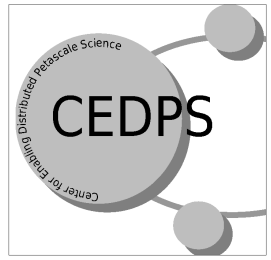


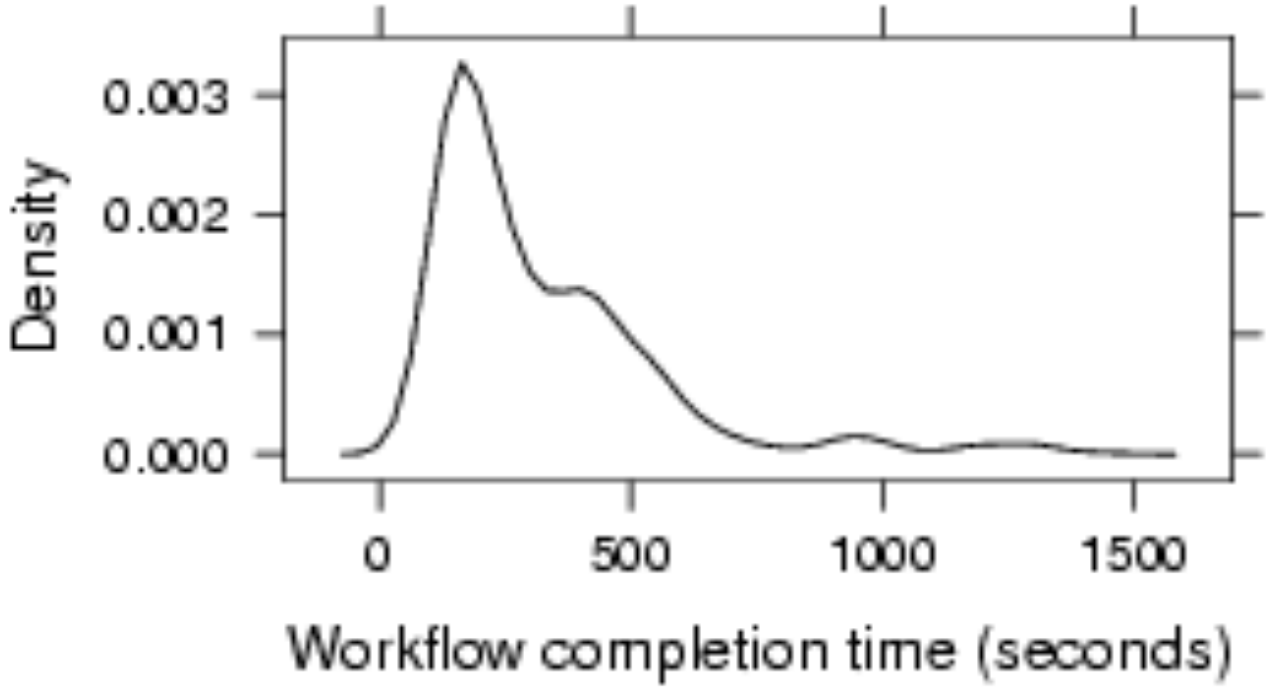


Q: What is the cause of the very long tail?

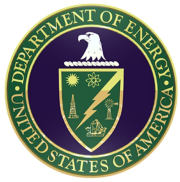


SciDAC
Scientific Discovery through Advanced Computing

# Distributed Debugging Issues

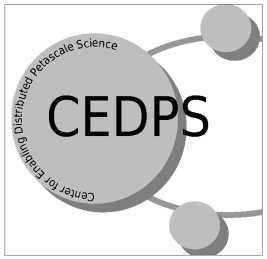


- Multiple Log file formats
- Inconsistent logging levels
- What resources / hosts are being used?
  - Grid's partially hide this
- Naïve Solution
  - Log on to each host that is part of your Grid job (if possible)
  - locate the log files
  - grep for error messages
- Very tedious process!!


SciDAC
Scientific Discovery through Advanced Computing

# Key Troubleshooting Components

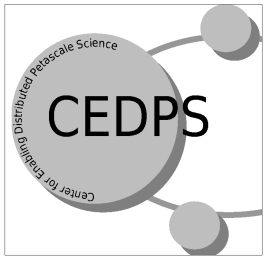


- Instrumented applications and middleware
  - Preferably with a common log format
- Time synchronized hosts
  - Run NTP everywhere (www.ntp.org)
- Log collection tools
- Log analysis tools

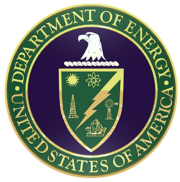


SciDAC
Scientific Discovery through
Advanced Computing

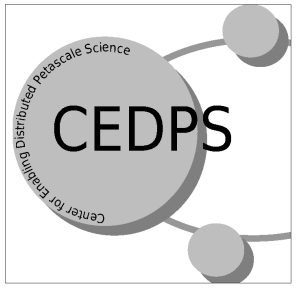


# Unified Logging: A "Best Practices" Guide for Grid Log Messages

Brian L. Tierney, Daniel Gunter: LBNL

Jen Schopf: ANL

Laura Perlman: ISI

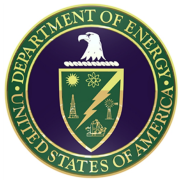


SciDAC
Scientific Discovery through Advanced Computing

# A Case for Unified Logging

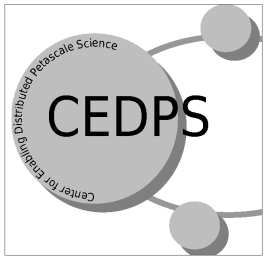


- Unified logging is needed for the Grid
  - Troubleshooting
  - Auditing
  - Forensics
- Example use cases:
  - My job failed. Why? Need to look at error logs on several hosts and several sites.
  - What is the list of hosts/resources that a given compromised user credential accessed in the past week


SciDAC
Scientific Discovery through
Advanced Computing

# Unified Logging

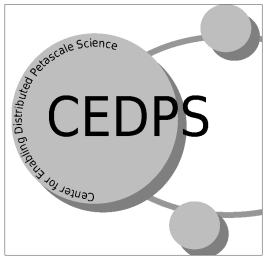


- Unified logging consists of 2 components
  - "Standard" log format
    - and converters for legacy logs
  - Log file collection mechanism

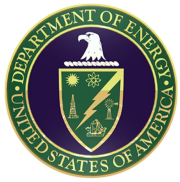


SciDAC
Scientific Discovery through Advanced Computing

# Log Collection

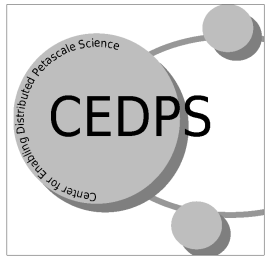


- No need to invent something new for this
  - syslog-ng fills all requirements
    - Open source, runs on all major OSes
    - Fault tolerant, secure (via stunnel), scalable, easy to configure, etc.
    - Large user base

- http://www.balabit.com/products/syslog-ng/

SciDAC
Scientific Discovery through Advanced Computing

# Log collection using syslog-ng

CEDPS
Center for Enabling Distributed Petascale Science

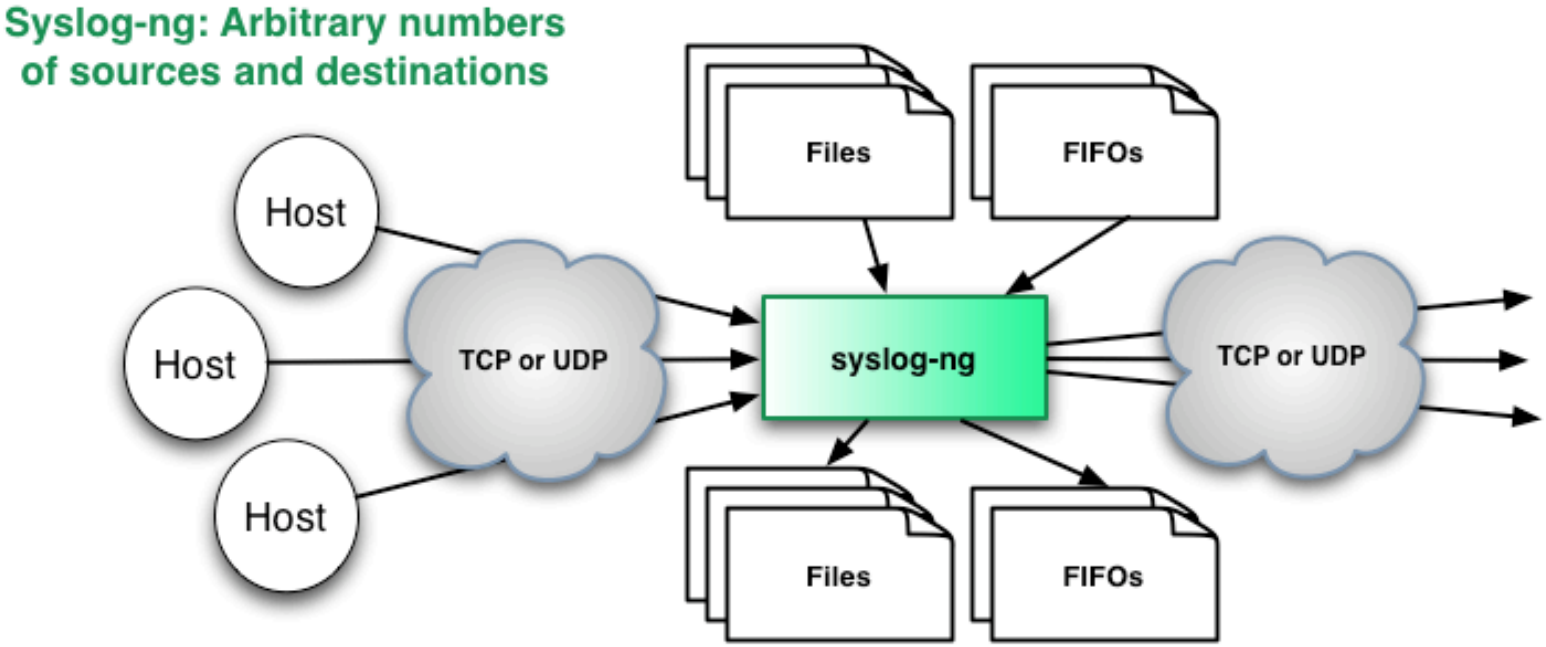

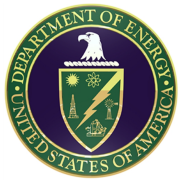


SciDAC
Scientific Discovery through Advanced Computing

# syslog-ng Features

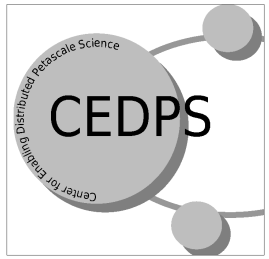


- Features:
  - Can filter logs based on level and **content**
  - Arbitrary number of sources and destinations
  - Provides remote logging
    - Can act as a proxy, tunnel thru firewalls
  - Execute programs
    - Send email, load database, etc.
  - Built-in log rotation
  - Timezone support
  - Fully qualified host names
  - Secure via stunnel (http://www.stunnel.org)
    - allows you to encrypt arbitrary TCP connections inside SSL

SciDAC
Scientific Discovery through Advanced Computing

# Sample Site Deployment

CEDPS

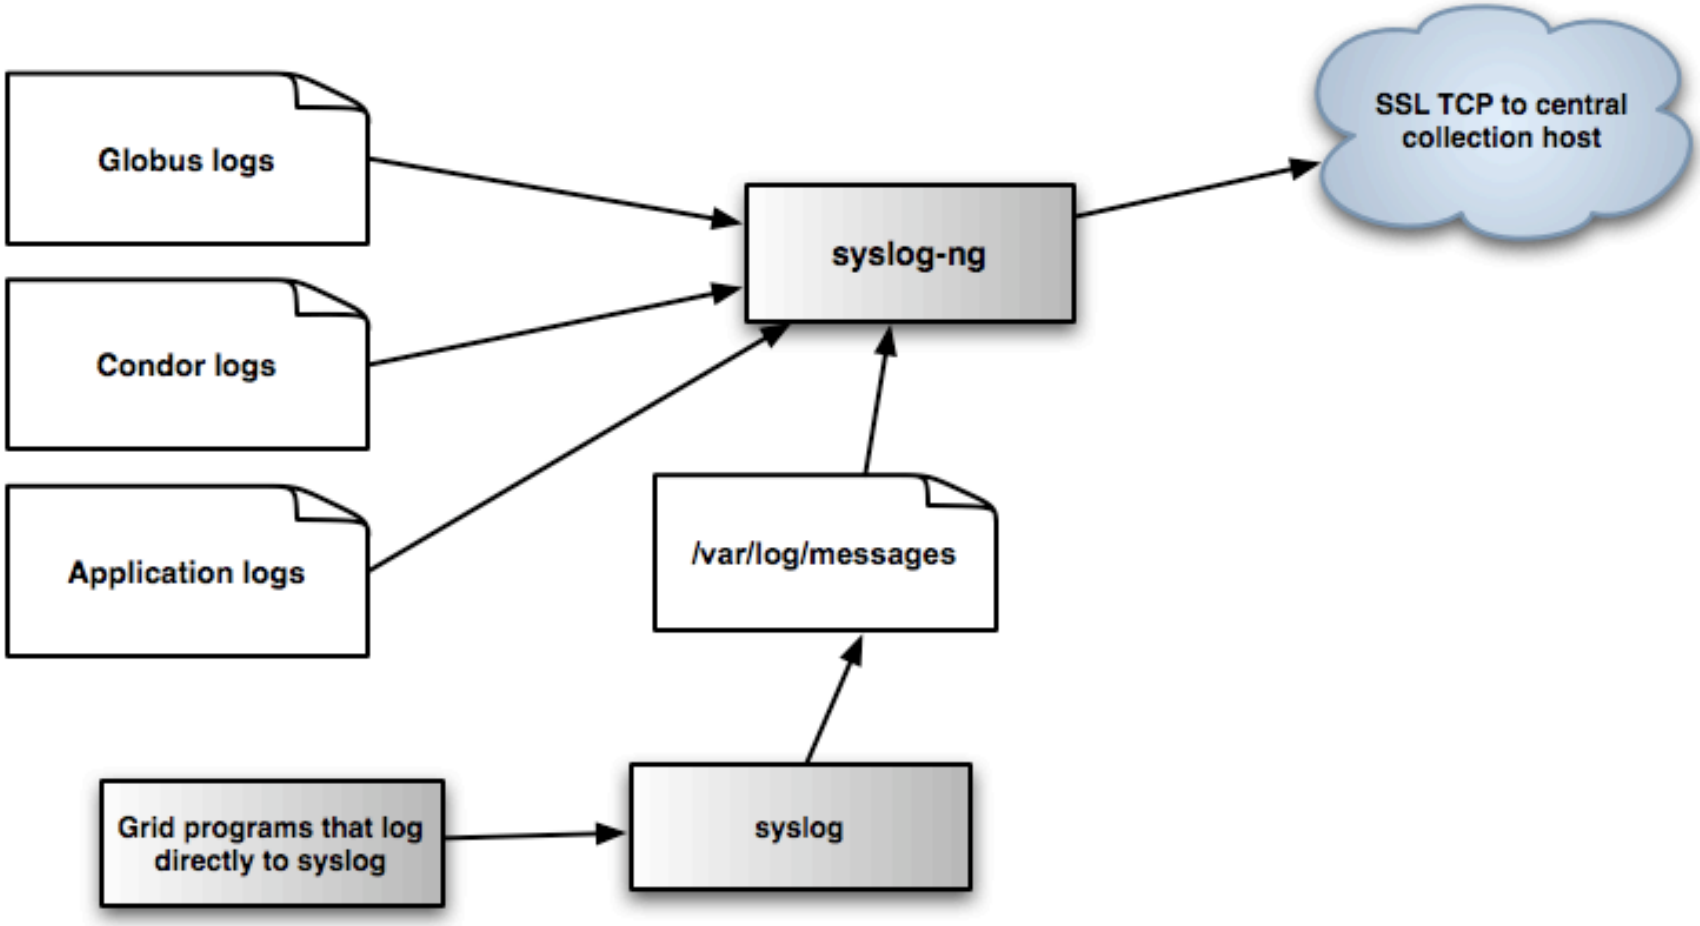

# Syslog-ng Deployment for OSG

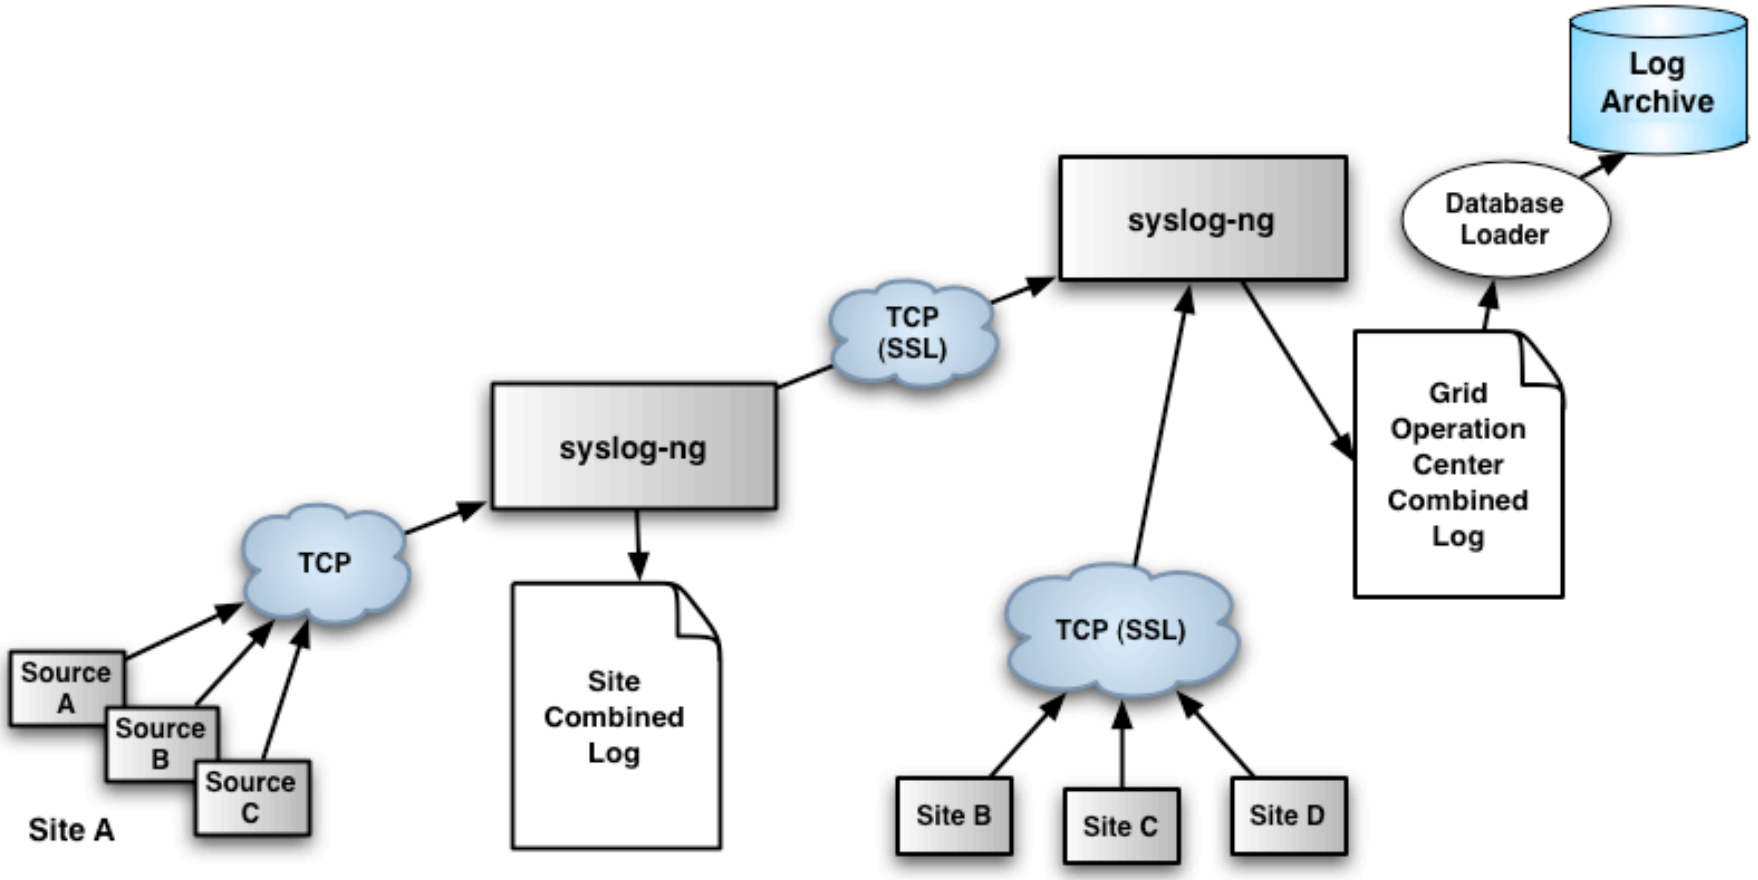


SciDAC
Scientific Discovery through Advanced Computing

# Logging "Best Practices" Recommendations

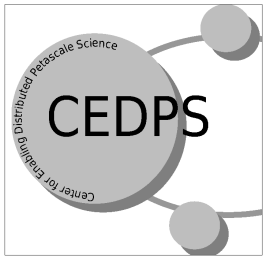


- Practices
  - All logs should contain a unique event name and an ISO-format timestamp
  - All system operations that might fail or experience performance variations should be wrapped with start and end events.
  - All logs from a given execution thread should be tagged with a globally unique ID (or GUID), such as a Universal Unique Identifiers (UUIDs)
- Log format
  - Logs should be composed of lines of ASCII name=value pairs
  - Example:ts=2006-12-08T18:48:27.598448Z
    event=org.globus.gridFTP.transfer.start prog=GridFTP-v4.2
    guid=1DDF1F3D-A677-4DBC-8C4E-6A8A3B252AE3
    file=filename src.host=H1 src.port=P1 dst.host=H2 dst.port=P2

http://www.cedps.net/wiki/index.php/LoggingBestPractice

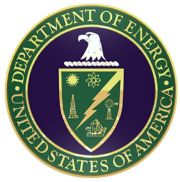


SciDAC
Scientific Discovery through Advanced Computing

# Event Names

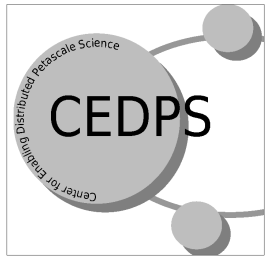


- Use a '.' as a separator and go from general to specific
  - Same as Java class names
- First part of name should be used as a unique namespace (e.g.: org.globus)
- Use start/end suffixes whenever possible
  - Helps immensely with troubleshooting
- Examples
  - `org.globus.gridFTP.start`
  - `org.globus.gridFTP.authn.start`
  - `org.globus.gridFTP.authn.end`
  - `org.globus.gridFTP.transfer.start`
  - `org.globus.gridFTP.transfer.end`
  - `org.globus.gridFTP.end`

```
-org.globus.MDS.response.start
-org.globus.MDS.query.start
-org.globus.MDS.query.end
-org.globus.MDS.write.net.start
-org.globus.MDS.write.net.end
-org.globus.MDS.response.end
```


SciDAC
Scientific Discovery through Advanced Computing

# Globally Unique IDs

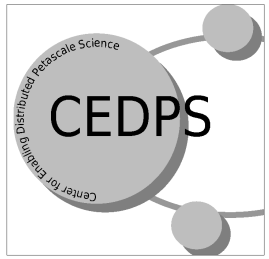


- Use the 'guid' or 'id' reserved name to allow correlation of a set of events together
  - `event=org.globus.gridFTP.authn.start id=27023`
  - `event=org.globus.gridFTP.authn.end id=27023`
  - `event=org.globus.gridFTP.transfer.start id=27023`
  - `event=org.globus.gridFTP.transfer.end id=27023`
- Can use standard unix/windows program 'uuidgen' to generate globally unique ID
  - e.g.: A5A563CD-D80C-4E58-9ECD-79C6B611E122

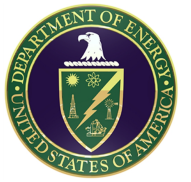


SciDAC
Scientific Discovery through Advanced Computing

# Reporting Errors

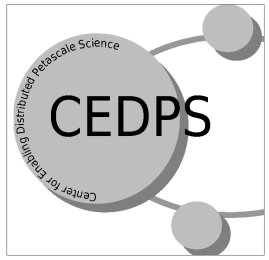


- Errors should be reported as part of the 'end' event if possible
  - Use 'status=N' (>= 0 success)
  - Not attempting to define other status codes
    - too hard to get agreement on these
- Example:
  - ```
    ts=2006-12-08T18:39:23.114369Z
    event=org.globus.authz.gridmap.end
     status=-1 DN="/O=CEDS/CN=Some User"
     msg="Cannot open gridmap file /etc/grid-security/
    grid-mapfile for reading"
     guid=F7D64975-069A-4152-A21F-57109AA46DFA
     level=ERROR
    ```
    ```

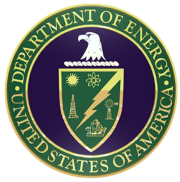


SciDAC
Scientific Discovery through Advanced Computing

# Ability to Filter "Sensitive" Data

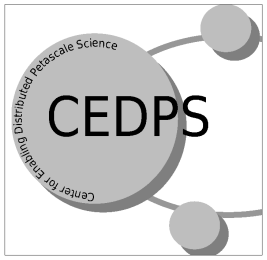


- Use the following reserved names make it possible to strip out sensitive data
  - DN – X509 distinguished name
  - user –user name
  - file – file name
  - dir – directory name
  - IP / localIP / remoteIP – local/remote IP address
  - host / localhost / remotehost – local/remote host name


SciDAC
Scientific Discovery through Advanced Computing

# How to Instrument Your Application

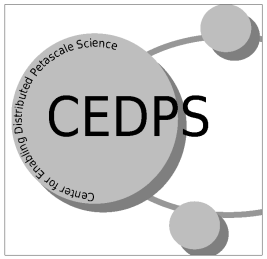


- You'll probably want to add instrumentation code to the following places in your distributed application:
  - before and after all disk I/O
  - before and after all network I/O
  - before and after any database query
  - entering and leaving each distributed component
  - before and after any significant computation
    - e.g.: an FFT operation
  - before and after any significant graphics call
    - e.g.: certain CPU intensive OpenGL calls

- This is usually an iterative process
  - add more instrumentation points as you zero in on the problem


SciDAC
Scientific Discovery through Advanced Computing

# Example: GridFTP

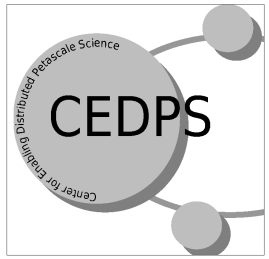


```
ts=2006-12-08T18:39:23.114369Z event=org.globus.gridFTP.start
   prog=GridFTP-4.0.3 localhost=myhost remoteHost=somehost.gov:56010
   serverMode=inetd guid=1DDF1F3D-A677-4DBC-8C4E-6A8A3B252AE3
ts=2006-12-08T18:39:23.114567Z event=org.globus.gridFTP.authn.start DN="/
   DC=org/DC=doegrids/OU=People/CN=Somebody" guid=1DDF1F3D-
   A677-4DBC-8C4E-6A8A3B252AE3
ts=2006-12-08T18:39:25.514369Z event=org.globus.gridFTP.authn.end DN="/
   DC=org/DC=doegrids/OU=People/CN=Somebody" msg="123456 successfully
   authorized" localUser=uscmspool381 guid=1DDF1F3D-
   A677-4DBC-8C4E-6A8A3B252AE3 status=0
ts=2006-12-08T18:39:25.864369Z event=org.globus.gridFTP.transfer.start
   file=/tmp/myfile tcpBufferSize=128KB dataBlockSize=262144 numStreams=1
   numStripes=1 destHost=129.79.4.64 guid=1DDF1F3D-
   A677-4DBC-8C4E-6A8A3B252AE3
ts=2006-12-08T18:45:02.214369Z event=org.globus.gridFTP.transfer.end file=/
   tmp/myfile bytesTransferred=678433 guid=1DDF1F3D-
   A677-4DBC-8C4E-6A8A3B252AE3 status=0
ts=2006-12-08T18:45:02.214386Z event=org.globus.gridFTP.end guid=1DDF1F3D-
   A677-4DBC-8C4E-6A8A3B252AE3 status=226
```

SciDAC
Scientific Discovery through
Advanced Computing

# CEDPS Troubleshooting Architecture

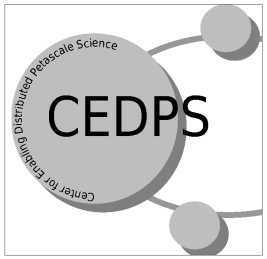


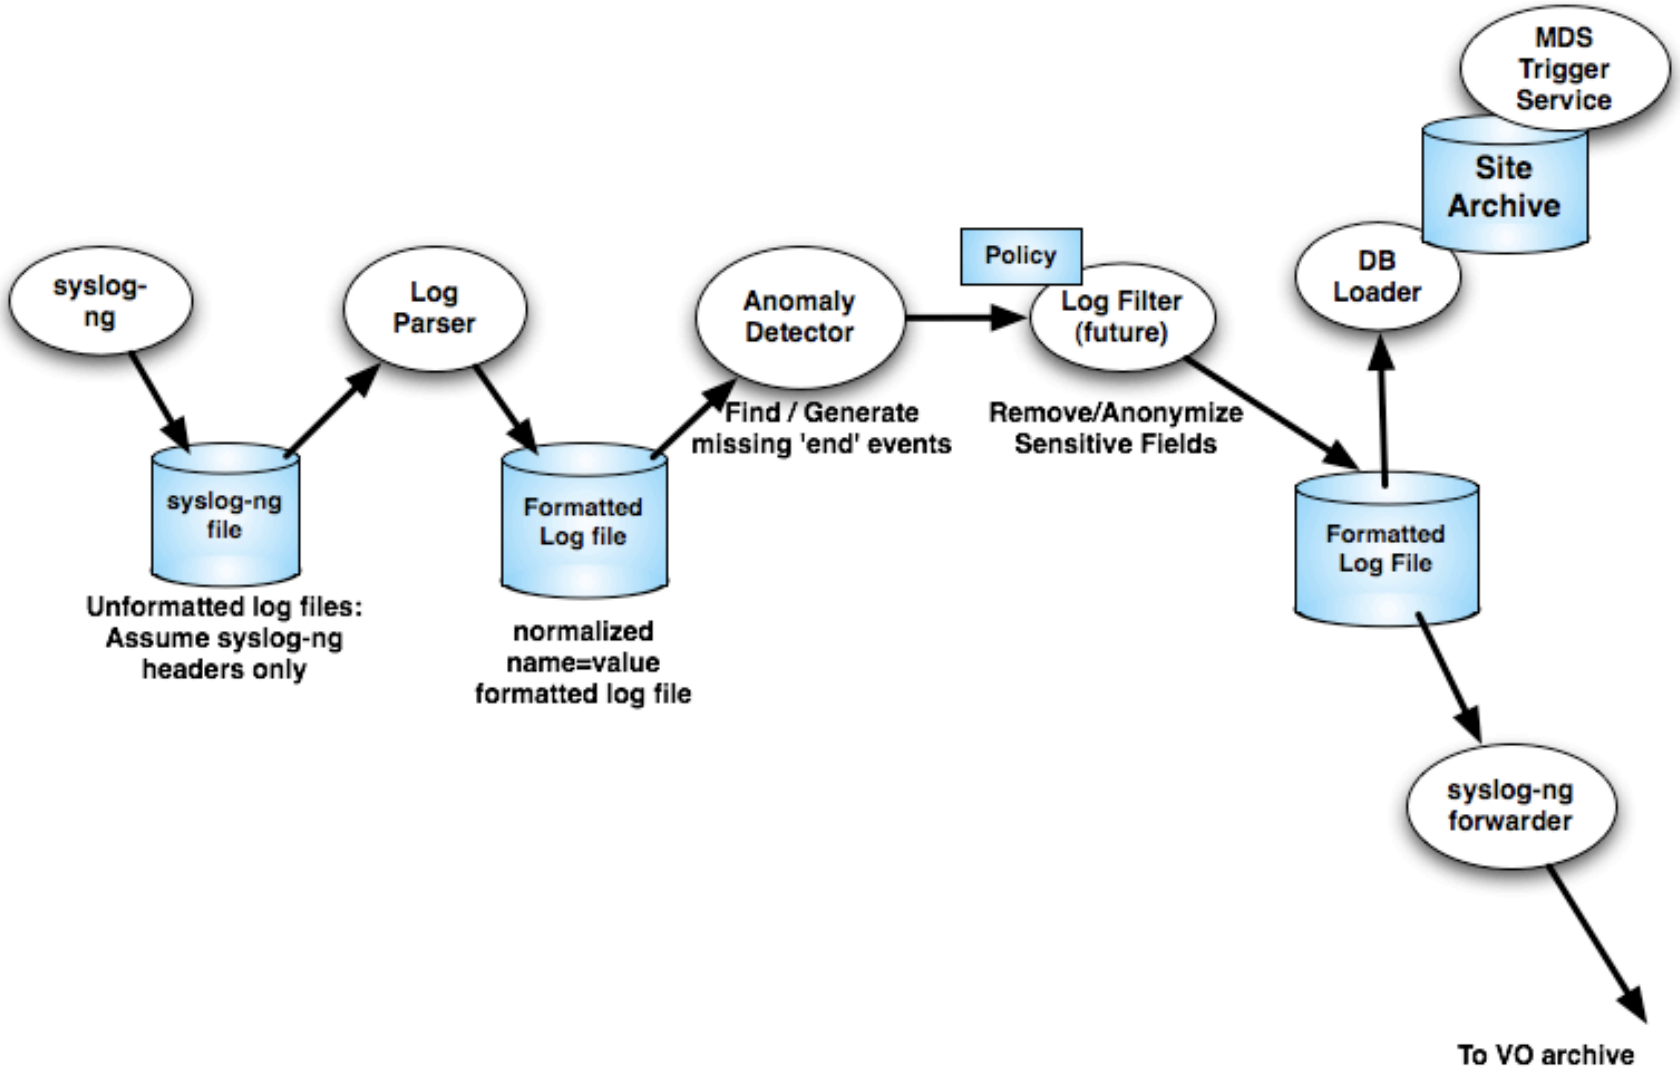


# CEDPS Application Partners

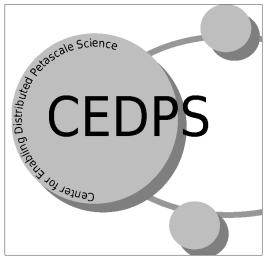


- Globus:
  - currently changing to "best practice" format logs
  - will build MDS "triggers" based on log database
- OSG:
  - Working closely with OSG on syslog-ng and stunnel configuration and deployment strategies
- Star project:
  - will use "best practice" format
  - working with them on GUID propagation issues, log database, alerts, etc.

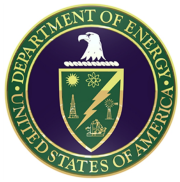


SciDAC
Scientific Discovery through
Advanced Computing

# More Information

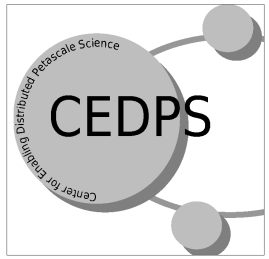


- http://www.cedps.net/wiki/index.php/Troubleshooting
- email: BLTierney@lbl.gov