

# bring network and time together using Linux tracing

Alexander Aring

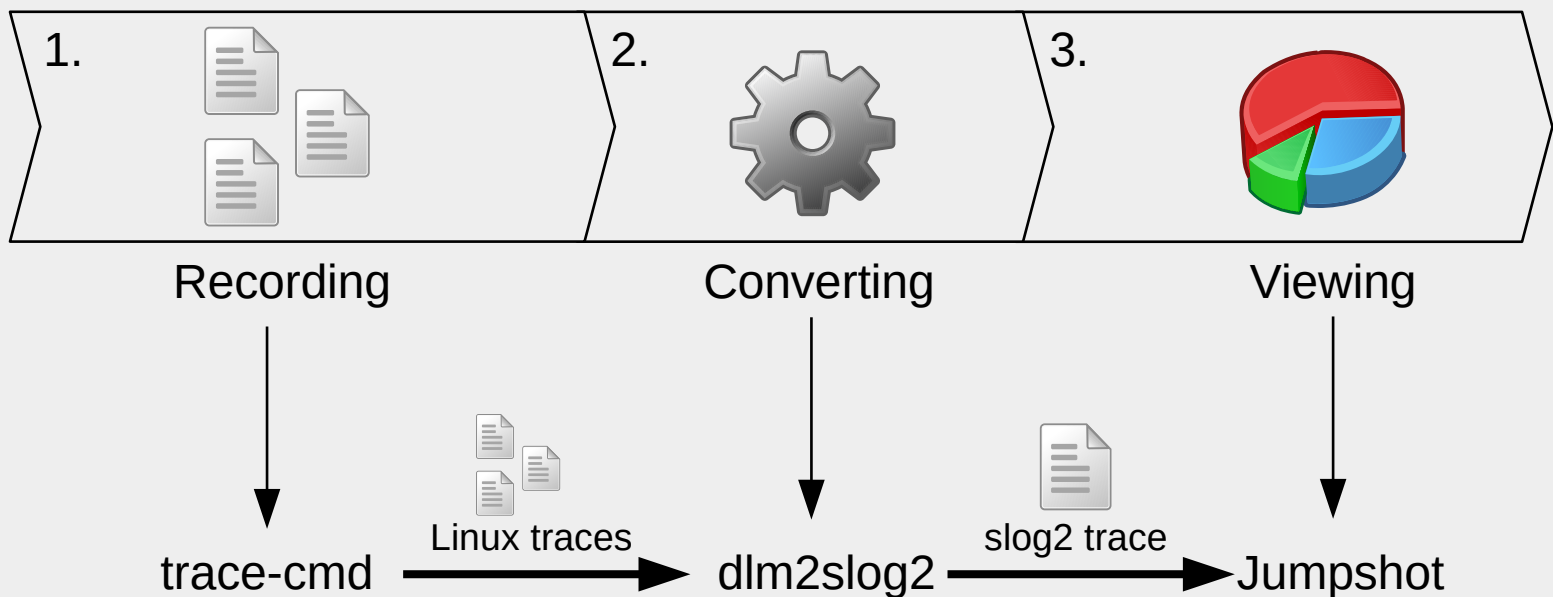
---

# Software Components

---


- trace-cmd (GPL/LGPL)
- slog2sdk (MIT)
  - NOT REQUIRED
  - Written in Java for MPI Apps by LANS
  - Library to create slog2 trace files
  - Viewer Jumpshot for slog2 files
- dlm2slog2 (BSD-3)
  - Trace Converter (Bridge Linux - Slog2)
  - Developed in this Project (for DLM)

# Our Pipeline Steps




# Visualizing Traces (slog2)


## Events



## States



## Example (GANTT)




## Legend

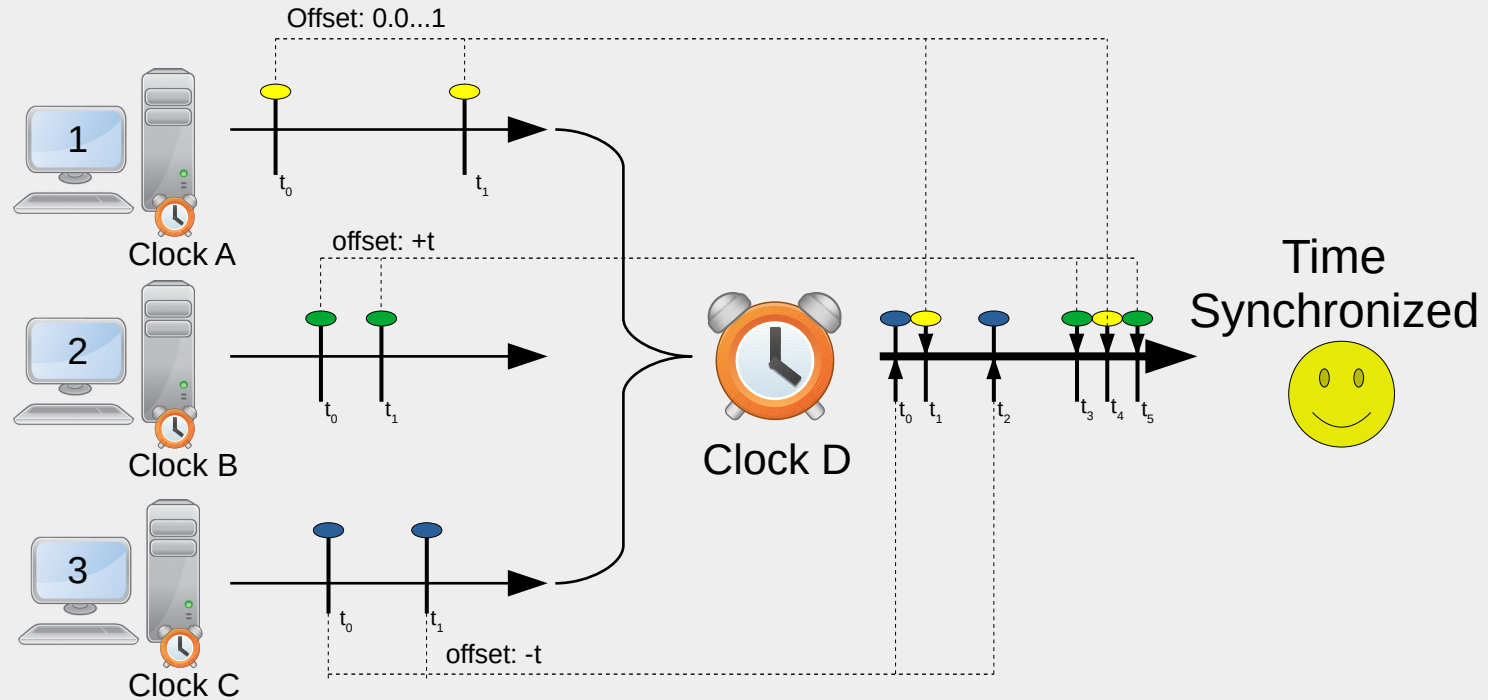
- Event1
- Event2
- Event3
- Event4
- State0

# Local Tracing

- Clocks are in reality more complex
- Clocks are shown as trace-cmd user



# Time synchronized Tracing



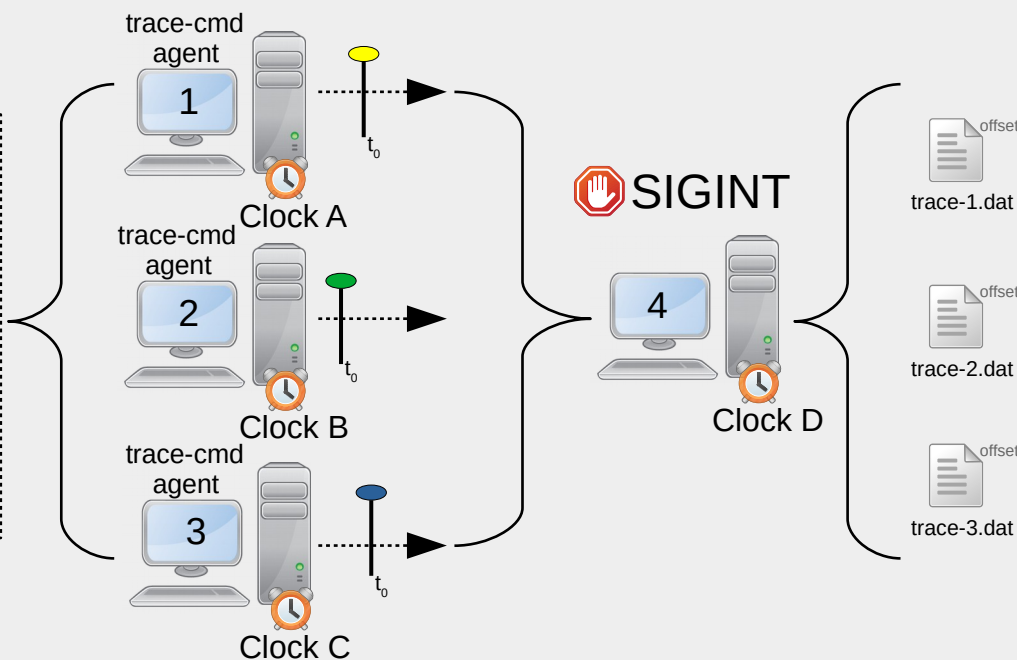
# trace-cmd Time Synchronization



```
trace-cmd record \  
-A 1 -e ... -A 2 -e ...
```



- Filtering
- TimeSync
  - KVM
  - PTP
- Socket
  - IP
  - VSOCK



# Next Slides Notice!

---

We close the Time Synchronization chapter


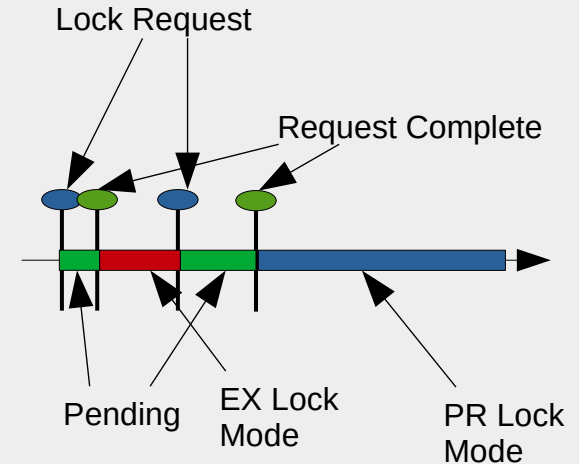
Every Timeline shown in the next slides are  
Time Synchronized!

We will look now how we can use it for DLM!

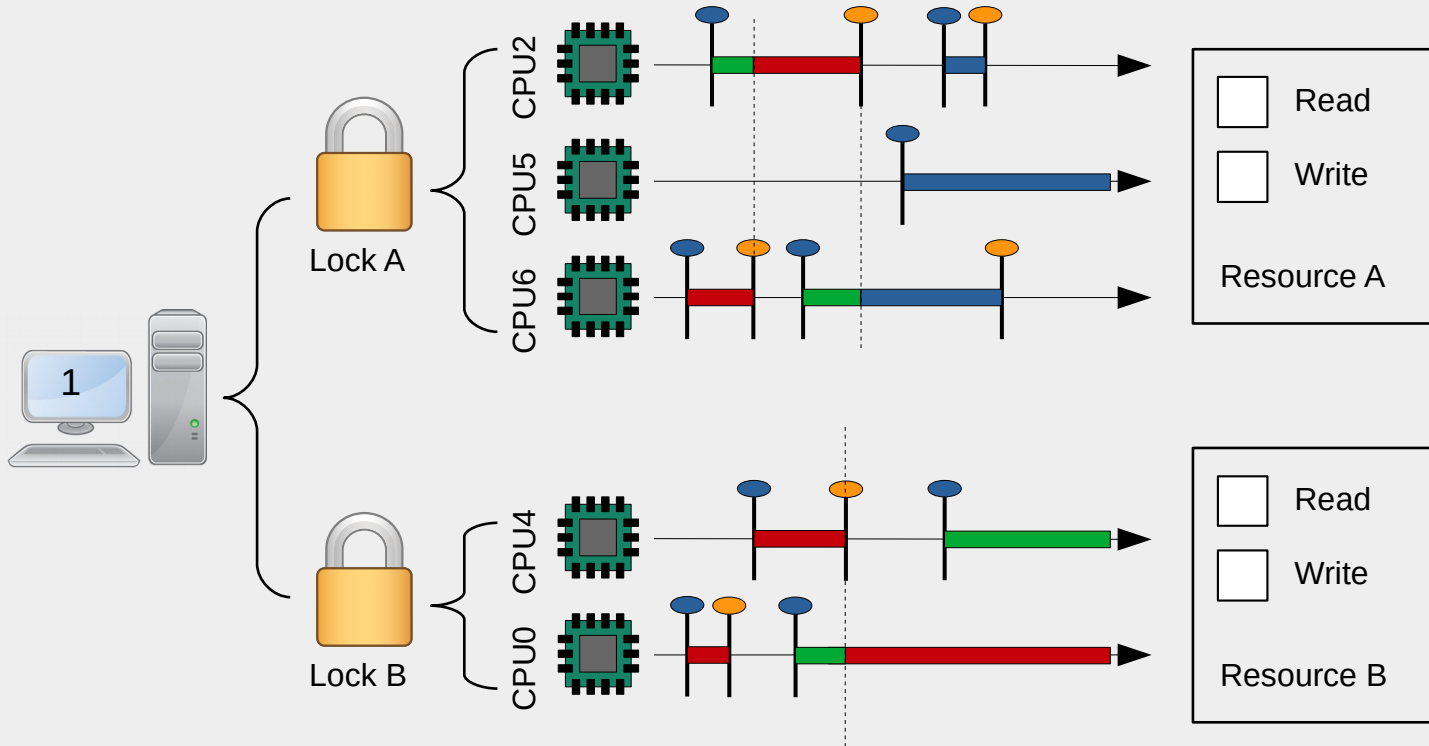


# Linux Distributed Lock Manager

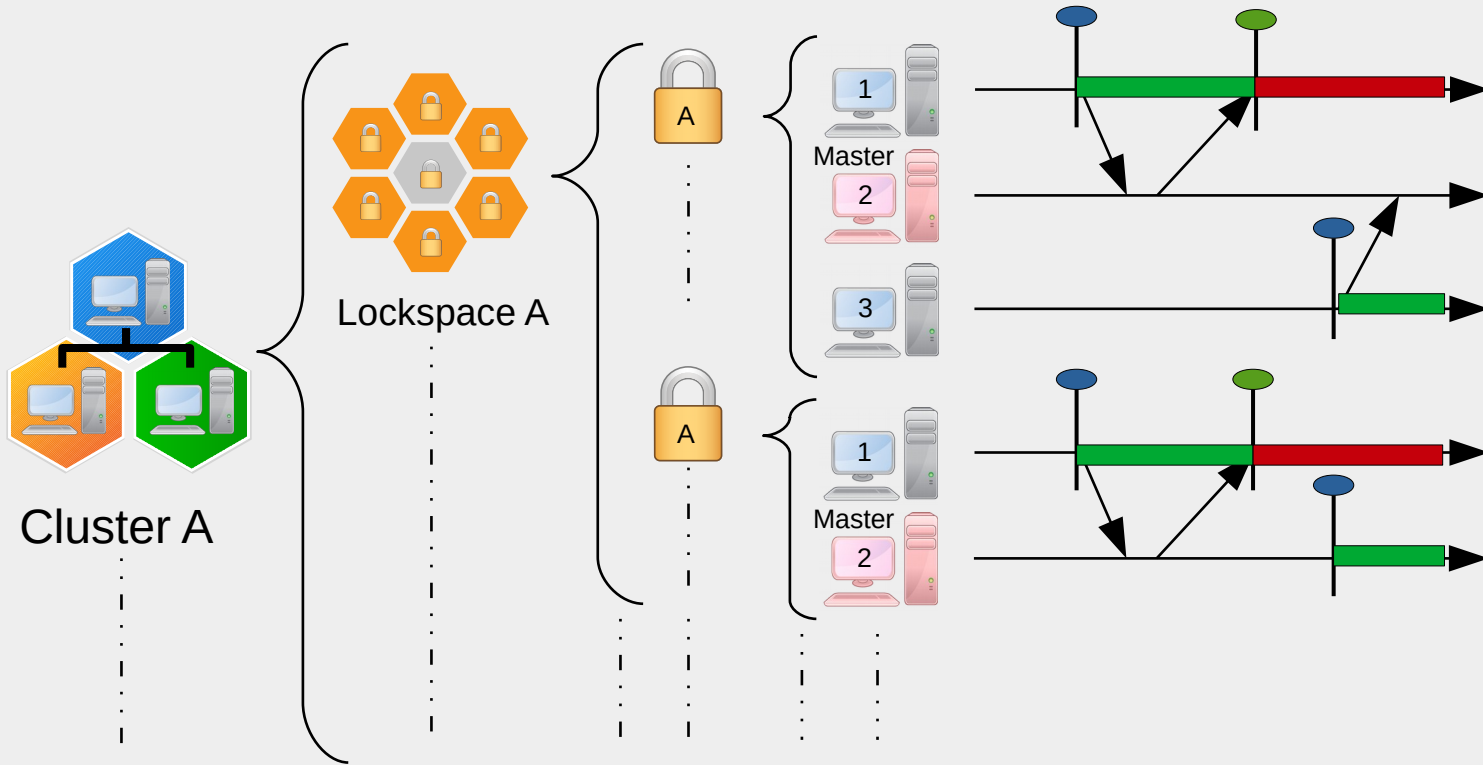
- Networking Protocol
- Requires Cluster Manager
- Asynchronous API
- Locks have Lock Modes
- Principle Lock Masters



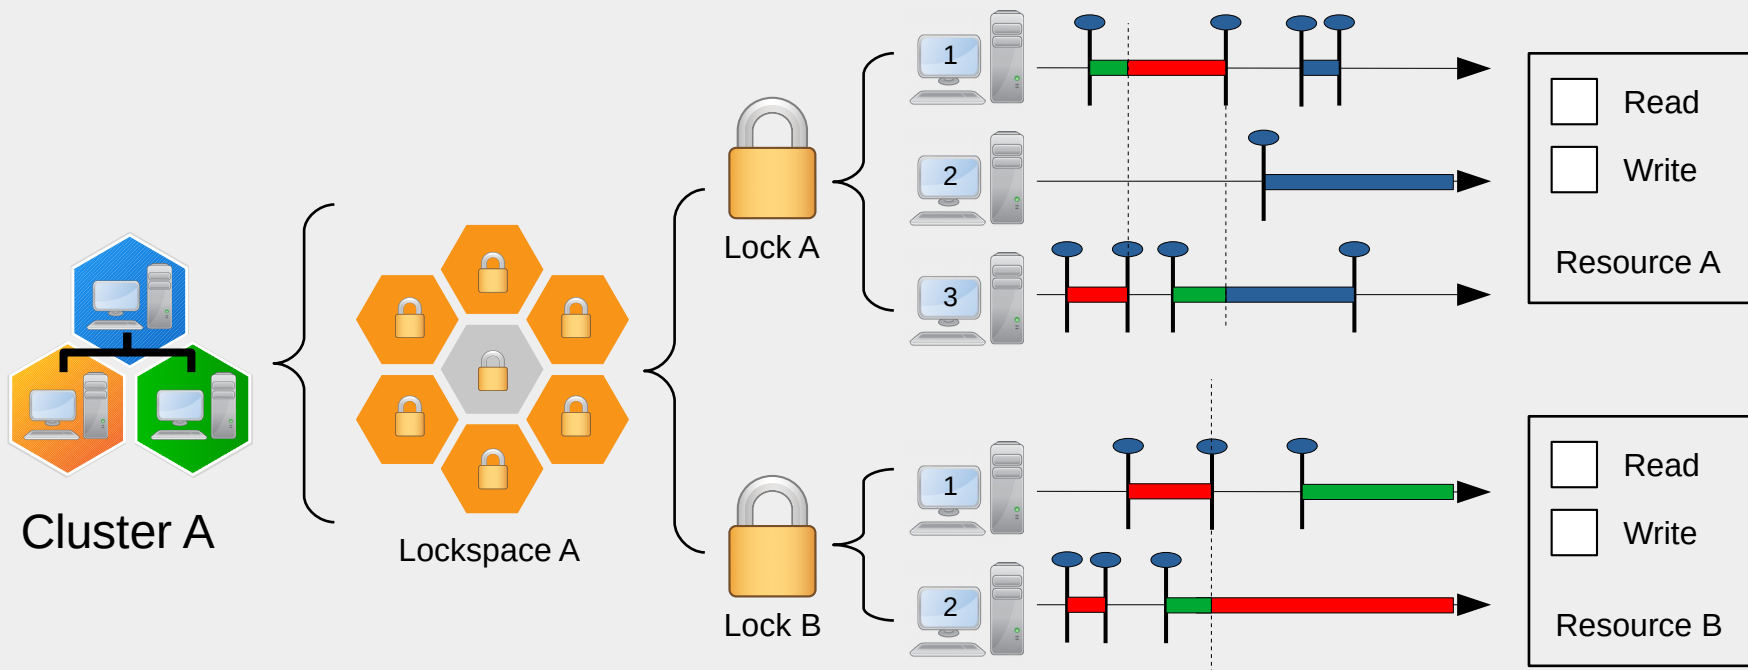
# Traced Local Locking



# DLM Linux-Cluster Hierarchy

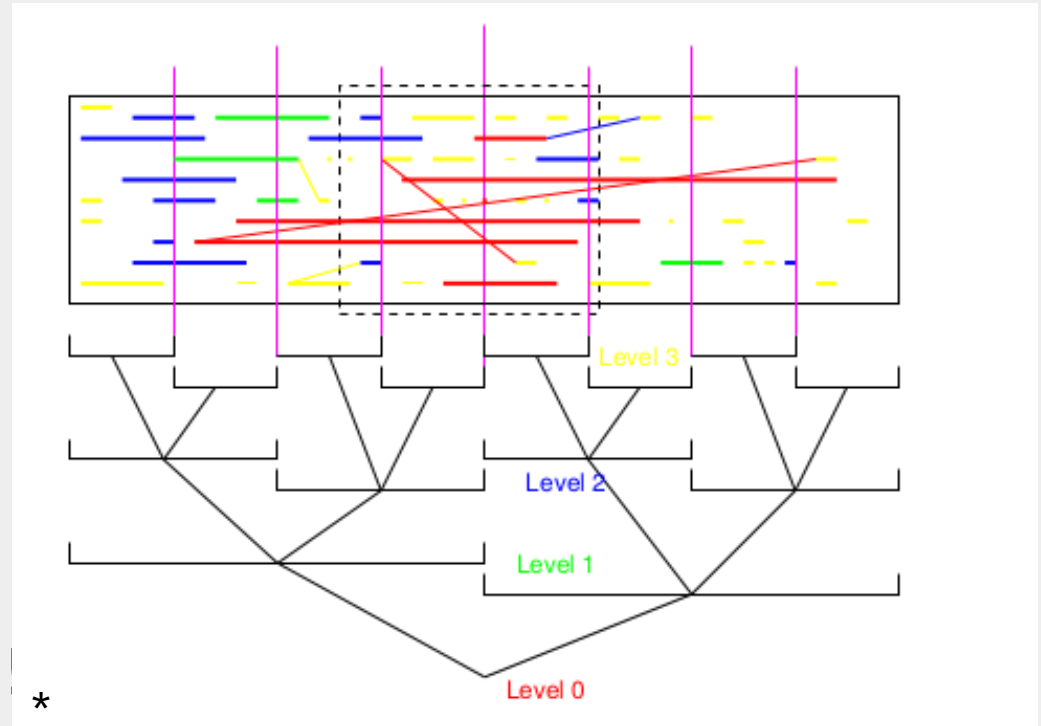


# Traced Clustered DLM Locking (compatible modes)



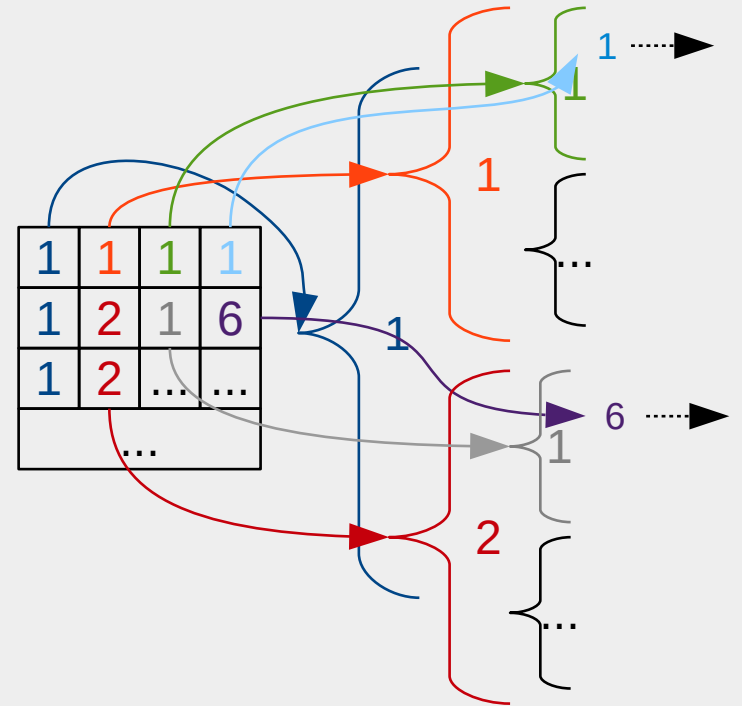
# SLOG2 Format Timeline Scaling

- Long Trace Record
- Storage is Cheap
- Scaling?
  - Level Of Detail
  - Tree-based
  - Partial-Read
  - Preview Events
  - X-Axis Scale
  - Y-Axis may not scale



# SLOG2 YCoordMap

- Multidimensional Array
- Columns for Y-Axis
- Column to LineID\*
- DLM Cluster Hierarchy



# dIm2slog2

## 1) Analyze

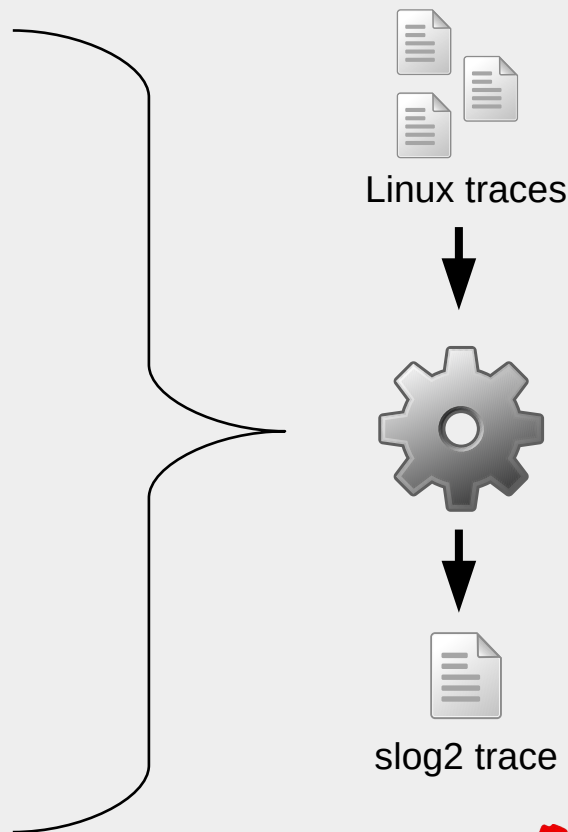
- Java-Bindings (Parsing)
- Available IDs, Locks, etc.

## 2) Filter

- Provide Way Scale Y-Axis
- Requires User-Interaction

## 3) Generate

- Events to Lock Modes States
- YCoordMap



# dlm2slog2



Menu About

Add Node

- Nodeid: 1	Select Linux Trace Filepath	./examples/dlm_traces/trace-5.dat
- Nodeid: 2	Select Linux Trace Filepath	./examples/dlm_traces/trace-4.dat
- Nodeid: 3	Select Linux Trace Filepath	./examples/dlm_traces/trace-3.dat

Analyze Resources

Available Resources				Filtered Resources			
#Events	Lockspace ID	Resource Name		#Events	Lockspace ID	Resource Name	
315	0x08e4f48a18	3	1d9d9f	1551	0x08e4f48a18	2	450ea7
315	0x08e4f48a18	3	805d	1063	0x08e4f48a18	3	449eb3
308	0x08e4f48a18	3	11	370	0x08e4f48a18	5	451252
266	0x08e4f48a18	2	45107a	364	0x08e4f48a18	3	4037
263	0x08e4f48a18	3	91120f	327	0x08e4f48a18	3	1f9b37
246	0x08e4f48a18	5	451121	299	0x08e4f48a18	5	45107a
182	0x08e4f48a18	2	451252	222	0x08e4f48a18	3	6f3af7
180	0x08e4f48a18	5	45108e				
134	0x08e4f48a18	2	451121				
123	0x08e4f48a18	5	45107e				
111	0x08e4f48a18	5	45107b				
105	0x08e4f48a18	5	45107c				
84	0x08e4f48a18	2	45107c				
81	0x08e4f48a18	5	451088				
78	0x08e4f48a18	2	45107e				
71	0x08e4f48a18	2	45107b				
70	0x08e4f48a18	5	200557				
67	0x08e4f48a18	5	2005cf				
67	0x08e4f48a18	5	917b5f				
67	0x08e4f48a18	5	917b5e				
66	0x08e4f48a18	5	917b6a				
66	0x08e4f48a18	5	2005db				
62	0x08e4f48a18	5	45107d				
61	0x08e4f48a18	2	917b6a				
61	0x08e4f48a18	2	4513e8				
60	0x08e4f48a18	5	4514e3				
60	0x08e4f48a18	5	45112b				
58	0x08e4f48a18	2	26c46b				
58	0x08e4f48a18	2	451088				
58	0x08e4f48a18	2	200557				
58	0x08e4f48a18	2	917b5e				

Select Slog2 Filepath: ./example.slog2

Generate and Save

100%


Linux Traces

Available

Filtered

To Slog2





# Jumpshot

Legend

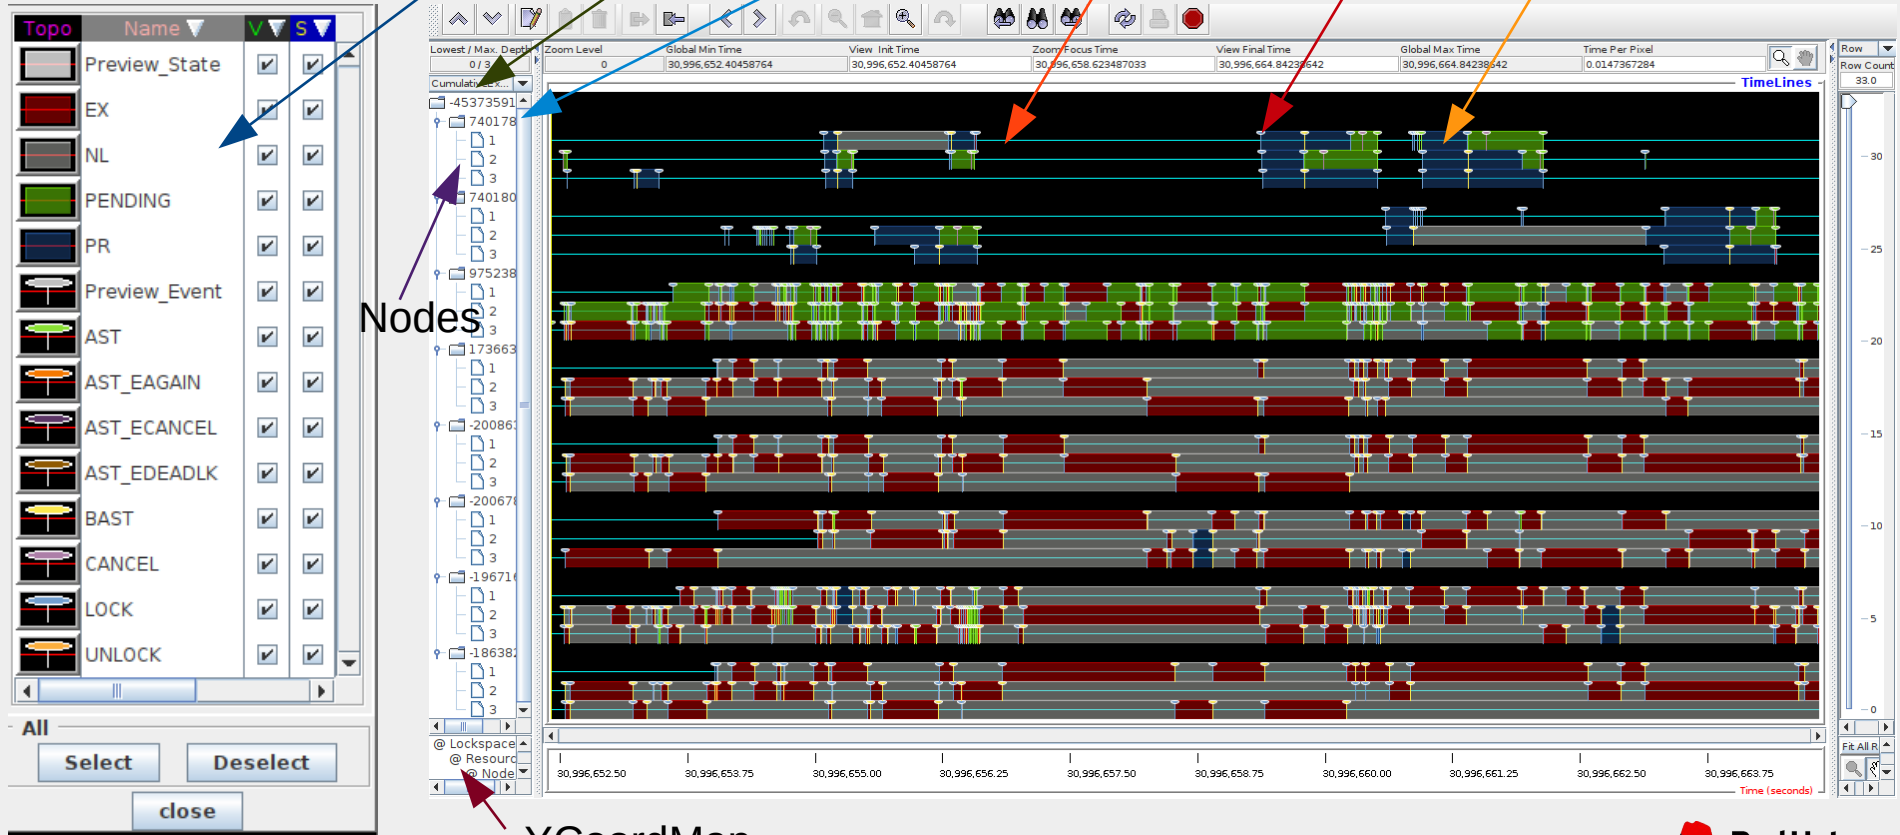
Lock

Lockspace

Timelines

Events

States



# Future Work

---

- Continuous Kernel Integration
  - DLM Locktorture Test
  - Check for Lock incompatible Modes
- Runtime Kernel Optimization
  - Predigt and Switch Lock Master
  - Networking Queues with skb mark?

**End**

---

Thanks