



Root Cause Analysis

October 15, 2008

LACSS Resilience Workshop

Presented by Jon Stearley

Sandia National Laboratories

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.



Agenda

(10 slides, 3 goals)

Get your help on Root Cause

- How to represent interdependencies?

Encourage you towards standardized validation data

- Component Operations Status (COS)

Make you aware of some other work

- Sisyphus (Logs)
- 9Lives (OS)

Context: Resilience Activities at Sandia

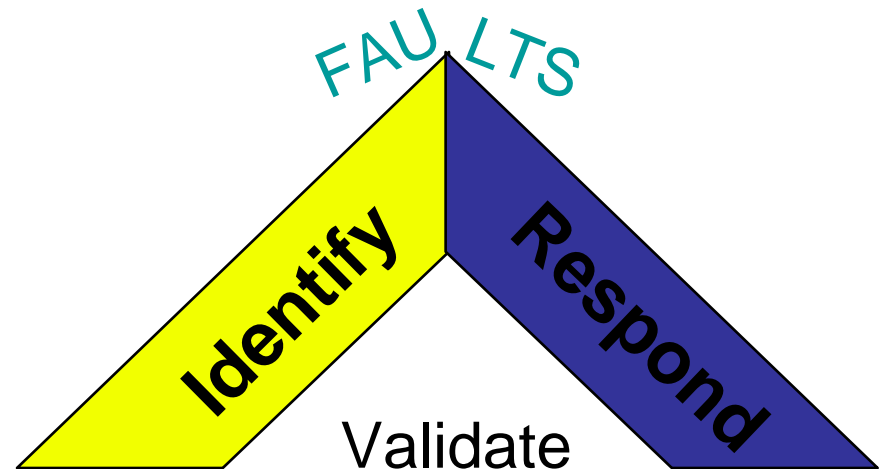
Goal: Automatic *identification* and *response* to faults

Identify Faults (CSSE)

- Failure Prediction
- Root cause analysis
- *Impact: enable timely and focused response*

Respond (LDRD):

- System-Directed Resilience
- *Impact: enable apps to run continuously despite faults*



Root Cause: Big Cheese Model

* Holes: incomplete data, incomplete time, incomplete components, incomplete dependencies

Time

Data

Components & Dependencies!

QuickTime™ and a
IFF (Uncompressed) decompressor
are needed to see this picture.



What Stinks?!#

Given hints (influences search path):

Distinguishing Symptoms

- Text (logs, username, job id, rank id, ...)
- Numbers (temperature, correlation, ...)

Suspect Components

- Hardware (racks, cores, routes, ...)
- Software (daemons, apps, libraries, ...)

Dependencies

- Functional, Physical
- Static, Dynamic

Identify:

Root Cause (likelihood ranked)

Root Cause: Big Cheese Model

* Holes: incomplete data, incomplete time, incomplete components, incomplete dependencies

Given hints (influences search path):

LABELS

Distinguishing Symptoms

- Text (logs, username, job id, rank id, ...)
- Numbers (temperature, correlation, ...)

VERTICES

Suspect Components

- Hardware (racks, cores, routes, ...)
- Software (daemons, apps, libraries, ...)

EDGES

Dependencies

- Functional, Physical
- Static, Dynamic

quickTime™ and a compressed) decomposed to see this picture

$G=(V,E,L)$

Properties?

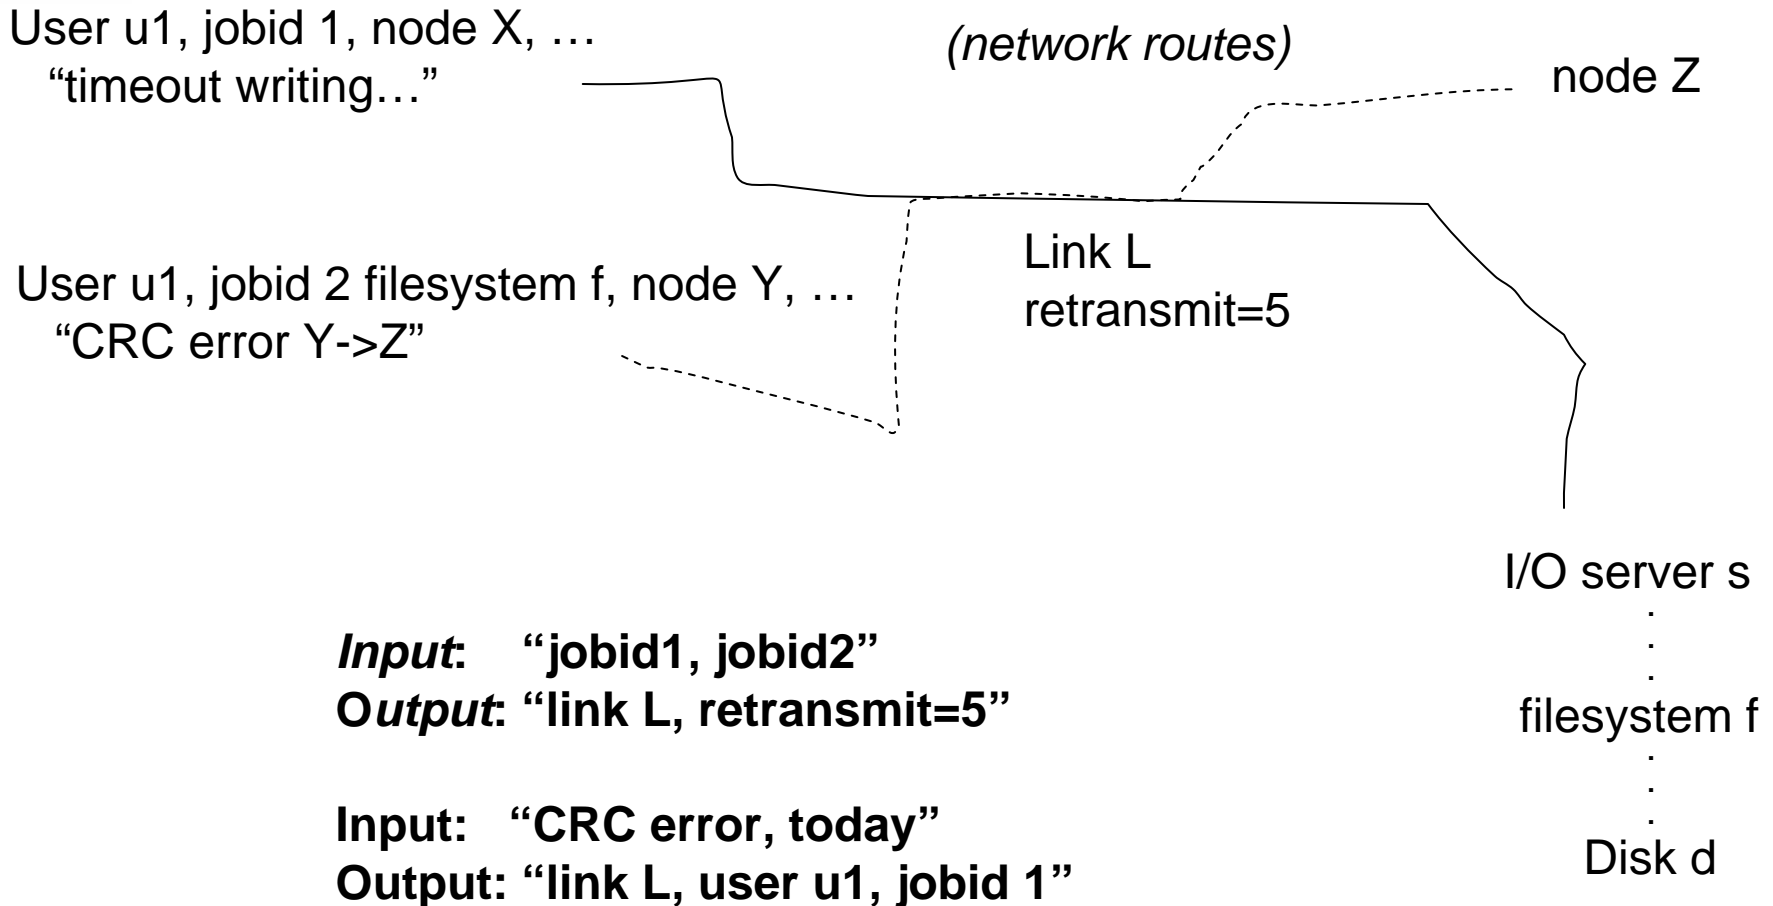
Useful model?

Useful algorithms?

Identify:

Root Cause (likelihood ranked)

Root Cause: Flakey Link (e.g.)



FY'09 Root Cause Deliverables

First Quarter

- Mathematical formulation of at least three important but currently non-computable root cause scenarios based on current systems

Second Quarter

- Select appropriate algorithms for solution.

Third Quarter

- Demonstrate proof-of-concept solutions.

Fourth Quarter

- SAND report

Validation:

“Quantify our ability to predict node failures on the TLCC platform”

Component Operations Status (COS)

Production Uptime (PU) = ready for immediate use by one or more production user

Scheduled Downtime (SD) = not in PU for scheduled reasons

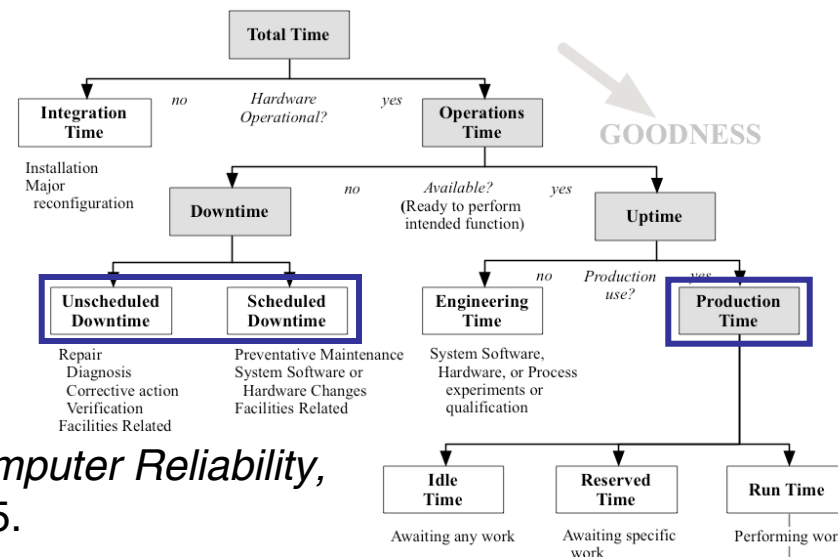
Unscheduled Downtime (UD) = not in PU for unscheduled reasons

FAILURE = the onset of Unscheduled Downtime

FY09:

1. Collect per-node COS on TLCC (admins express “SD” on CLI)
2. Perform prediction experiments on symptomatic data (numeric, text)
3. Use COS to quantify (and validate) prediction results

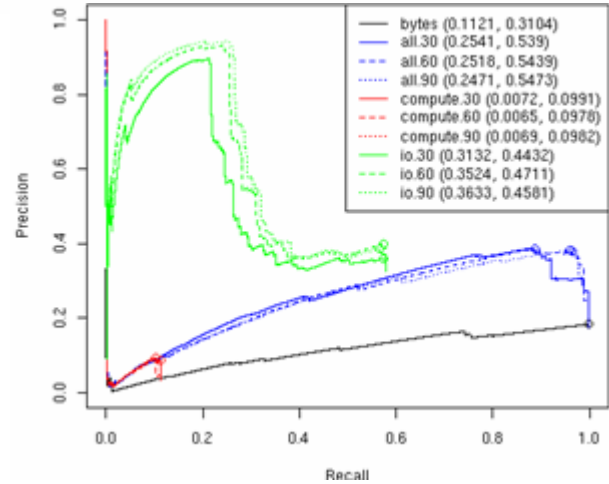


COS is a proposed standard



Stearley. *Defining and Measuring Supercomputer Reliability, Availability, and Serviceability (RAS)*, LCI05.

Recent Work on Logs (5 years worth)

Oliner, Stearley. *Alert Detection in System Logs*, ICDM'08.
www.cs.sandia.gov/sisyphus cfd.r.usenix.org/data.html#hpc4



bytes) [HELP] [change settings] Analyze:

```
Aug 4 14:00:12 10.1.0.84 local7 info DMT_EMT Unable to read disk 5S block 0x0002ED32 - Adding to Reassign List.
Aug 4 14:00:14 10.1.0.84 local7 info DMT_EMT Unable to read disk 5S block 0x0002ED56 - Adding to Reassign List.
Aug 4 14:00:15 10.1.0.84 local7 info DMT_EMT Unable to repair disk address 191794 LUN 8, 00000176 DLR:0, DLG:1, DRR:0, DEL:0, DELR:0
Aug 4 14:00:15 10.1.0.84 local7 info DMT_EMT Unable to repair disk address 191830 LUN 8, 00000176 DLR:0, DLG:1, DRR:0, DEL:0, DELR:0
Aug 4 14:00:15 10.1.0.84 local7 info DMT_EMT Unable to recover defects on disk 5S LUN 8, 00000176 DLR:0, DLG:1, DRR:0, DEL:0, DELR:0
Aug 4 14:00:15 10.1.0.84 local7 info DMT_EMT DMT EMT invalidate node LUN 8, 00000176 DLR:0, DLG:1, DRR:0, DEL:0, DELR:0, DERR:0
Aug 4 14:00:15 10.1.0.84 local7 crit DMT_EMT Warning: EMT multi-channel failure detection: BBM disabled: disk 5S, LUN 8, 00000176 DLR:0
DERR:0 r0 w0 ll fl0 fr2 ea:0,10
Aug 4 14:00:15 10.1.0.84 local7 info DMT_EMT Reload Node tier: 5 LUN 8, 00000176 DLR:0, DLG:1, DRR:0, DEL:0, DELR:0, DERR:0 r0 w0
Aug 4 14:00:17 10.1.0.84 local7 debug Disk_Int DISKS3: Received Rsp for incomplete exchange 00000021 Completed:00026400 Exp:00040000
Aug 4 14:00:17 10.1.0.84 local7 debug Disk_Int DISKS2: StartCmdRetry - Response for Drive 0x05S exch:00000021
Aug 4 14:00:17 10.1.0.84 local7 debug Disk_Int DISKS4: Cmd:28 LBA:0002EC00 Length:00040000
Aug 4 14:00:17 10.1.0.84 local7 debug Disk_Int DISKS4: Current Error SCSI Status:00000B02
Aug 4 14:00:17 10.1.0.84 local7 debug Disk_Int DISKS4: Under-run Residual Count:00019C00
Aug 4 14:00:17 10.1.0.84 local7 debug Disk_Int DISKS4: Response DL:00000008 Response:00000000 00000000
Aug 4 14:00:17 10.1.0.84 local7 debug Disk_Int DISKS4: Sense Key:03 Info:0002ED32
Aug 4 14:00:17 10.1.0.84 local7 debug Disk_Int DISKS4: ASC:11 ASCQ:00 FRU:E4
```

Responding to Faults (9Lives)

System-Directed Resilience (LDRD)

Three Primary Research Thrusts

- **Application quiescence**
 - Suspend App, handle in-flight messages and in-progress IO
- **Efficient state management**
 - Identify critical state (app characterization, app guidance, changesets)
 - System guidance for when to extract state (eg MTTI)
 - Explore diskless methods and compression
- **Fault recovery**
 - System software to support dynamic node allocation
 - Explore network virtualization to abstract physical node ID from software.
 - Explore efficient methods for state recovery (roll-back, roll-forward techniques)

Low Overhead (easier to develop and support, faster to run)

- Lightweight Kernel (Kitten)
- Stateless Networking Protocol (Portals)

End

Get your help on Root Cause

- How to represent interdependencies?

Convince you towards standardized validation metrics

- Component Operations Status (COS)

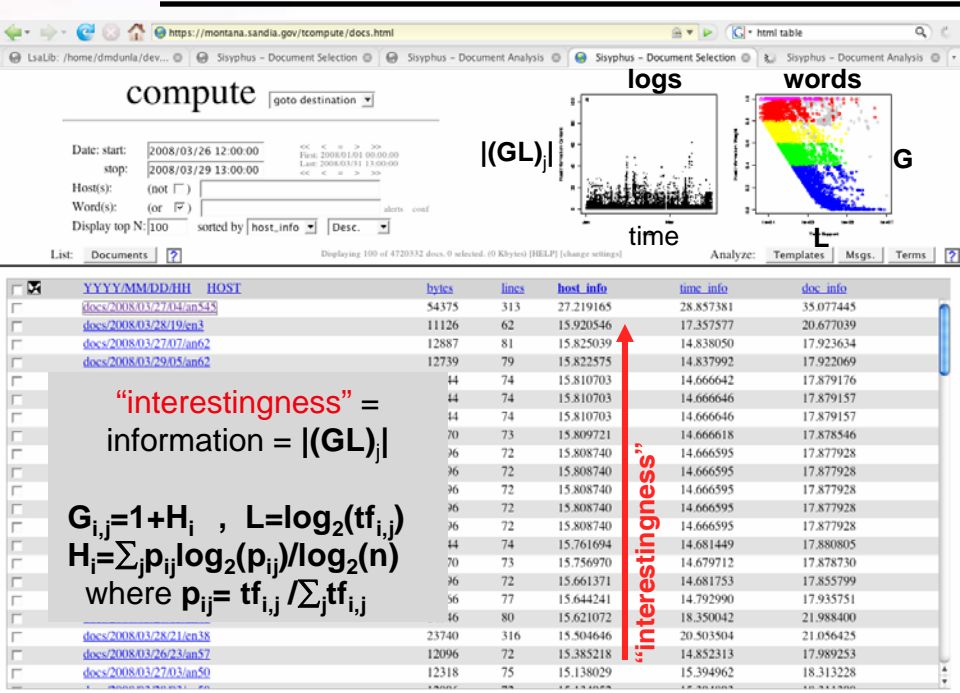
Make you aware of some other work

- Sisyphus (Logs)
- 9Lives (OS)

Extra Slides

Sisyphus

Automatic Fault Detection



compute [goto destination]

Date: start: 2008/03/26 12:00:00 stop: 2008/03/29 13:00:00

Host(s): (not) Words(s): (or) Display top N: 100 sorted by host_info Desc.

logs: $|(\text{GL})_j|$ vs time

words: Heatmap of word counts

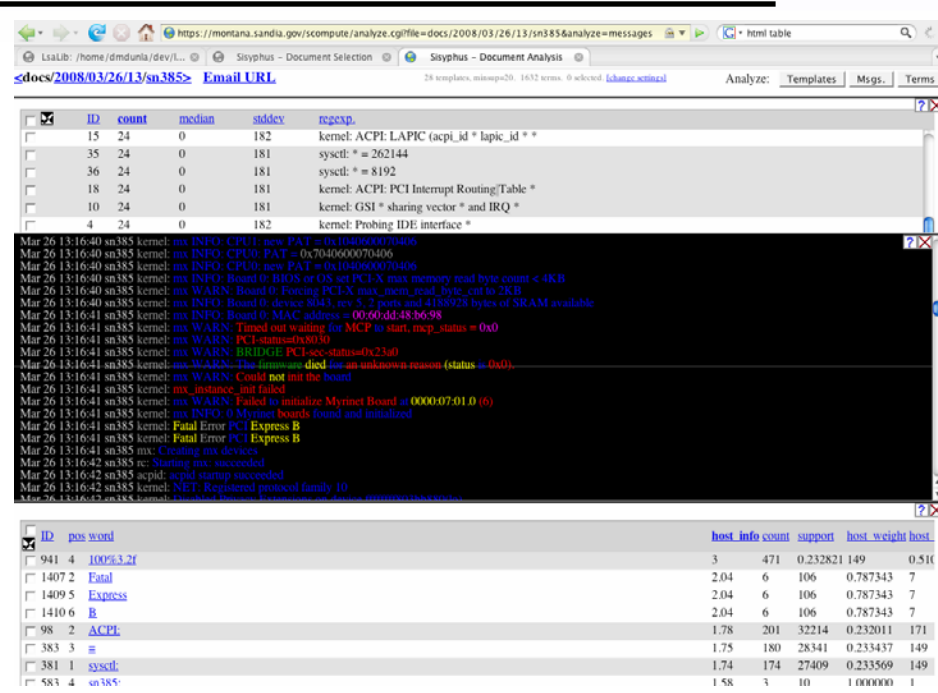
“interestingness” = information = $|(\text{GL})_j|$

$$G_{i,j} = 1 + H_i, \quad L = \log_2(\text{tf}_{i,j})$$

$$H_i = \sum_j p_{ij} \log_2(p_{ij} / \log_2(n))$$

where $p_{ij} = \text{tf}_{i,j} / \sum_j \text{tf}_{i,j}$

YYYYMMDD:HH	HOST	bytes	lines	host_info	time_info	doc_info
docs/2008/03/27/01/an549		54375	313	27.219165	28.857381	35.077445
docs/2008/03/28/01/an1		11126	62	15.920546	17.357577	20.677039
docs/2008/03/27/07/an62		12887	81	15.825039	14.838050	17.923634
docs/2008/03/29/05/an62		12739	79	15.822575	14.837992	17.922069



docs/2008/03/26/13/sn385: Email URL

ID	count	median	stddev	reexp.
15	24	0	182	kernel: ACPI: LAPIC (acpi_id * lapic_id * *
35	24	0	181	sysctl: * = 262144
36	24	0	181	sysctl: * = 8192
18	24	0	181	kernel: ACPI: Interrupt Routing/Table *
10	24	0	181	kernel: GSI * sharing vector * and IRQ *
4	24	0	182	kernel: Probing IDE interface *

Mar 26 13:16:40 sn385 kernel: mx: INFO: CPU0: new PAT = 0x1040600070406
 Mar 26 13:16:40 sn385 kernel: mx: INFO: CPU0: PAT = 0x7040600070406
 Mar 26 13:16:40 sn385 kernel: mx: INFO: CPU0: new PAT = 0x1040600070406
 Mar 26 13:16:40 sn385 kernel: mx: INFO: Board 0: BIOS or OS set PCI-A max memory read byte count < 4KB
 Mar 26 13:16:40 sn385 kernel: mx: WARN: Board 0: Forcing PCI-A max mem read byte cnt to 2KB
 Mar 26 13:16:40 sn385 kernel: mx: INFO: Board 0: device 8043, rev 5, 2 ports and 4188928 bytes of SRAM available
 Mar 26 13:16:41 sn385 kernel: mx: INFO: Board 0: MAC address = 00:60:dd:48:b6:98
 Mar 26 13:16:41 sn385 kernel: mx: WARN: Timed out waiting for MCP to start, mcp_status = 0x0
 Mar 26 13:16:41 sn385 kernel: mx: WARN: PCI-stmstatus=0x0
 Mar 26 13:16:41 sn385 kernel: mx: WARN: BRIDGE: PCI-sec-status=0x230f
 Mar 26 13:16:41 sn385 kernel: mx: WARN: The fuses are died for an unknown reason (status is 0x0)
 Mar 26 13:16:41 sn385 kernel: mx: WARN: Could not init the board
 Mar 26 13:16:41 sn385 kernel: mx: WARNING: init failed
 Mar 26 13:16:41 sn385 kernel: mx: WARN: Failed to initialize Myrinet Board at 0000:07:01:0 (6)
 Mar 26 13:16:41 sn385 kernel: mx: INFO: 0 Myrinet boards found and initialized
 Mar 26 13:16:41 sn385 kernel: Fatal Error PCI Express B
 Mar 26 13:16:41 sn385 kernel: Fatal Error PCI Express B
 Mar 26 13:16:41 sn385 mx: Creating mx device
 Mar 26 13:16:42 sn385 rc: Starting mx succeeded
 Mar 26 13:16:42 sn385 acpid: acpid startup succeeded
 Mar 26 13:16:42 sn385 kernel: NET: Registered protocol family 10
 Mar 26 13:16:42 sn385 kernel: Disabled Ethernet Controller on pci bus 0000:07:01:0

ID	pos	word	host_info	count	support	host_weight	host
941	4	100%3.2f		3	471	0.232821	149
1407	2	Fatal		2.04	6	106	0.787343
1409	5	Express		2.04	6	106	0.787343
1410	6	E		2.04	6	106	0.787343
98	2	ACPI:		1.78	201	32214	0.232011
383	3	=		1.75	180	28341	0.233437
381	1	sysctl:		1.74	174	27409	0.233569
583	4	sn385:		1.58	3	10	1.000000

1. Automatically rank logs by information content.

2. Automatically color words by information weight.

3. Automatically deduce word and message patterns.

Similar computers correctly performing similar work should produce similar logs (anomalies warrant investigation).

Identifying and Predicting Faults

Research to Identify Root Cause

- Develop mathematical methods to describe and track dependencies between system components
- Develop tools to analyze models, identify root cause, and provide feedback to fault-response systems (e.g., the system directed LDRD)

Research on Fault Prediction

- Develop methodologies for anomaly detection and quantification
- Correlate anomalous behaviors with root causes
- Use automated anomaly detection and learned correlations to predict failures with a calculated level of confidence

Impact

- Critical enabler for automated response (also useful to admins)
- Mathematics necessary to design and operate truly resilient systems
- Quantify the system-wide impact of failures.

FY'09 9Lives Deliverables

First Quarter

- Enumerate list of faults, likelihood of fault, and possible response (Stearley)
- Investigate options for quiescence (Brightwell)
- Investigate diskless state management and node recovery (Oldfield)
- Investigate related projects – BLCR (Kordenbrock)

Second Quarter

- Complete design of code to trace application memory usage (Pedretti)
- Complete design of network virtualization layer (Brightwell)
- Complete design of fake RAS system (Riesen)
- Complete design of diskless state management and recovery (Oldfield)

Fourth Quarter

- Memory-use characterizations for selected applications (Pedretti)
- Demonstrate response to link failures – may not need node recovery (Laros)
- Demonstrate network virtualization layer (Brightwell)
- Demonstrate fake RAS – inject faults (Riesen)

FY'09 Prediction Deliverables

- **First Quarter**
 - Selection of algorithms for anomaly detection (multivariate distributions and time series phenomena)
 - Define database schema for collection of system variables and for representation of data to support the different analyses
- **Second Quarter**
 - Quantification methodologies, failure definition and recording mechanisms, data collection (CPM, Log, Workload, COS, etc.)
- **Third Quarter**
 - Implementation and validation of anomaly detection algorithms
 - Selection of algorithms for correlation of failures with anomalies (classification and root cause inference)
 - Attributes selected for predictive analysis
 - Scalable data collection and representation
- **Fourth Quarter**
 - Implementation and validation of failure to anomaly correlation techniques
 - Write SAND report documenting:
 - Quantification of the ability to predict failures on the TLCC platform
 - Additional information/data that should be collected, as well as suggested scalable collection mechanisms, in order to improve failure predictability on future platforms

Proposed Process for Collecting per-node COS on TLCC

Scheduled
Downtime

Production
Uptime

Unscheduled
Downtime

SLURM and MOAB logs and databases are postprocessed into COS records.

All downtimes are counted as Unscheduled, unless distinguishable otherwise:

Only the admins know if a downtime is scheduled or not, so they must be involved in distinguishing them. The easiest way to do this would be to use a flag to scontrol and mrsvctl like:

```
scontrol -sd [cause...]  
mrsvctl -sd [cause...]
```

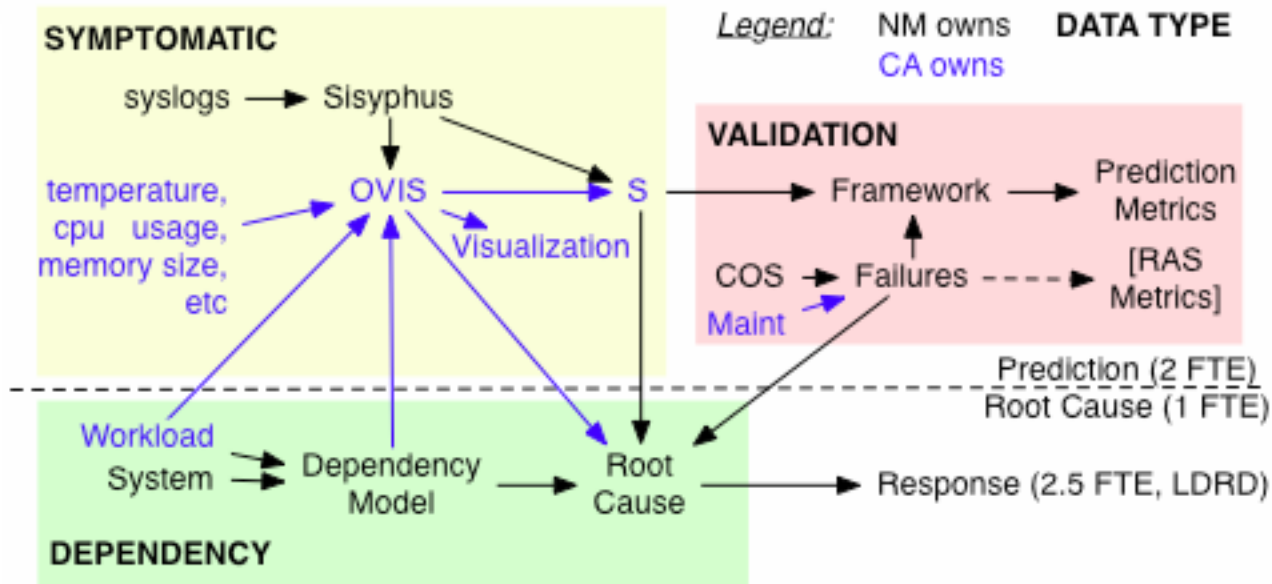
Which would be shorthand for

```
scontrol reason="SD: [cause]"  
mrsvctl -a ACCT=SD -D "SD: [cause]"
```

Where [cause] is an optional argument describing the downtime.

-sd could be added, aliased, or wrapped. More details are on the next slide.

SNL Resilience Interactions



Details:

Prediction Signal S: $s_i=(t,n,c)$ where t is time, n is node, and c is confidence

Failures F: $f_j=(t_i,n,d)$ where t_i is time of failure, n is node, and d is duration

COS is Component Operations Status (uptime, sched and unsched downtime)

Maint is maintenance records (hardware replacements etc)

Framework: how well does S predict F? $P(F | S)$

Prediction Metrics: Precision, Recall, AUC

Workload: jobs (user, nodes, duration,...) and executables (versions, libraries,...)

System: hardware and software configurations (cores, cables, routes, daemons,...)

Dependency Model: $G=(V,E)$ where V are components and E are dependencies

Root Cause: $P(G' | S,F)$ based primarily on dependency model

*Root Cause: $P(n | S)$ based primarily on symptomatic inference

[RAS Metrics]: no FY09 work is planned, but FYI COS is also key for:

MTTI, MTTR, component pareto (including CCF via Dependency Model), etc

Root Cause: Graph Model

* Holes: incomplete data, incomplete time, incomplete components, incomplete dependencies

Given hints (influences search path):

LABELS

Distinguishing Symptoms

- Text (logs, username, job id, rank id, ...)
- Thresholds (temperature limits, ...)
- Waveforms (cpu profile, ...)
- Correlations

VERTICES

Suspect Components **(paths)**

- Hardware (racks, cores, routes, ...)
- Software (daemons, apps, libraries, ...)

EDGES

Dependencies

- Functional, Physical
- Static, Dynamic

quickTime™ and a compressed) decomposed to see this picture

$G=(V,E,L)$

Properties?

Useful model?

Useful algorithms?

Identify:

Root Cause (likelihood ranked)