Elastic Stack Installation & Configuration

Anna Gassen, Ciara Goetze, James Gadson III
Team G Code

LLNL-PRES-XXXXXX
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

Objective

- Install and configure Elastic Stack on the Academy clusters
- Gather logs from all nodes
- Develop some insightful searches
- Research data analysis concepts

Elastic Stack

You know, for search

- Our clusters produce more than 1500 log messages per minute
- Comprised of six open-source tools: Elasticsearch, Logstash, Kibana, Beats, X-Pack, Elastic Cloud
- Allows quick analyzation, visualization, and mining of millions of log files
- Identify trends, statistics, and abnormalities

Logstash

- Collects data from many different sources at the same time
- Filters and parses each message, converts it into a common format for easier analysis
- Aggregates and transports data to Elasticsearch (or the software of your choice)

Filebeat

- A lightweight log file shipping agent
- Part of the Beats family of data shippers
- Communicates directly with Logstash or Elasticsearch
- Easily forwards and centralizes log files

Elasticsearch

- Full-text search engine that searches and centrally stores data
- Quickly find, retrieve, and analyze big volumes of data
- Distributed and highly scalable
- Near real time search
- Uses RESTful API, JSON, and Lucene

Kibana

- Data visualization tool for log and time series analytics
- Makes navigation and monitoring of logs more intuitive
- Provides numerous graph and dashboard options to display information

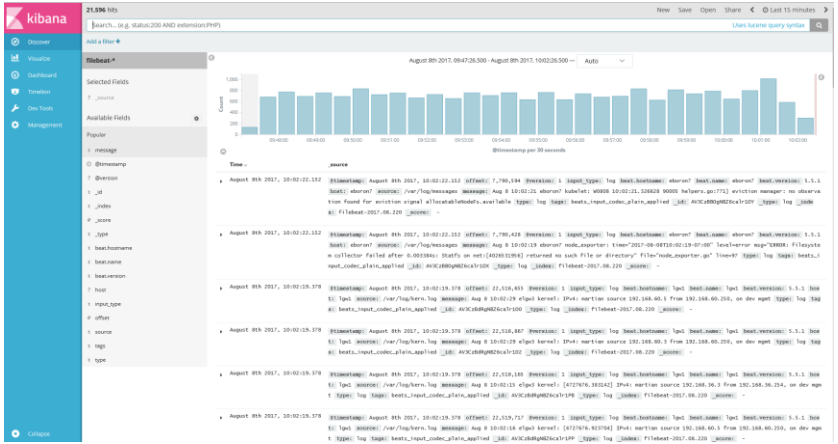

```
{
  "_index": "filebeat-2017.08.220",
  "_type": "log",
  "_id": "AV3CzBBOgNBZ6calr1OY",
  "_score": 1,
  "_source": {
    "@timestamp": "2017-08-08T17:02:22.152Z",
    "offset": 7790594,
    "@version": "1",
    "input_type": "log",
    "beat": {
      "hostname": "eboron7",
      "name": "eboron7",
      "version": "5.5.1"
    },
    "host": "eboron7",
    "source": "/var/log/messages",
    "message": "Aug  8 10:02:21 eboron7 kubelet: W0808 10:02:21.326628   90005 helpers.go:771] eviction manager: no observation found for eviction signal allocatableNodeFs.available",
    "type": "log",
    "tags": [
      "beats_input_codec_plain_applied"
    ]
  },
  "fields": {
    "@timestamp": [
      1502211742152
    ]
  }
}
```

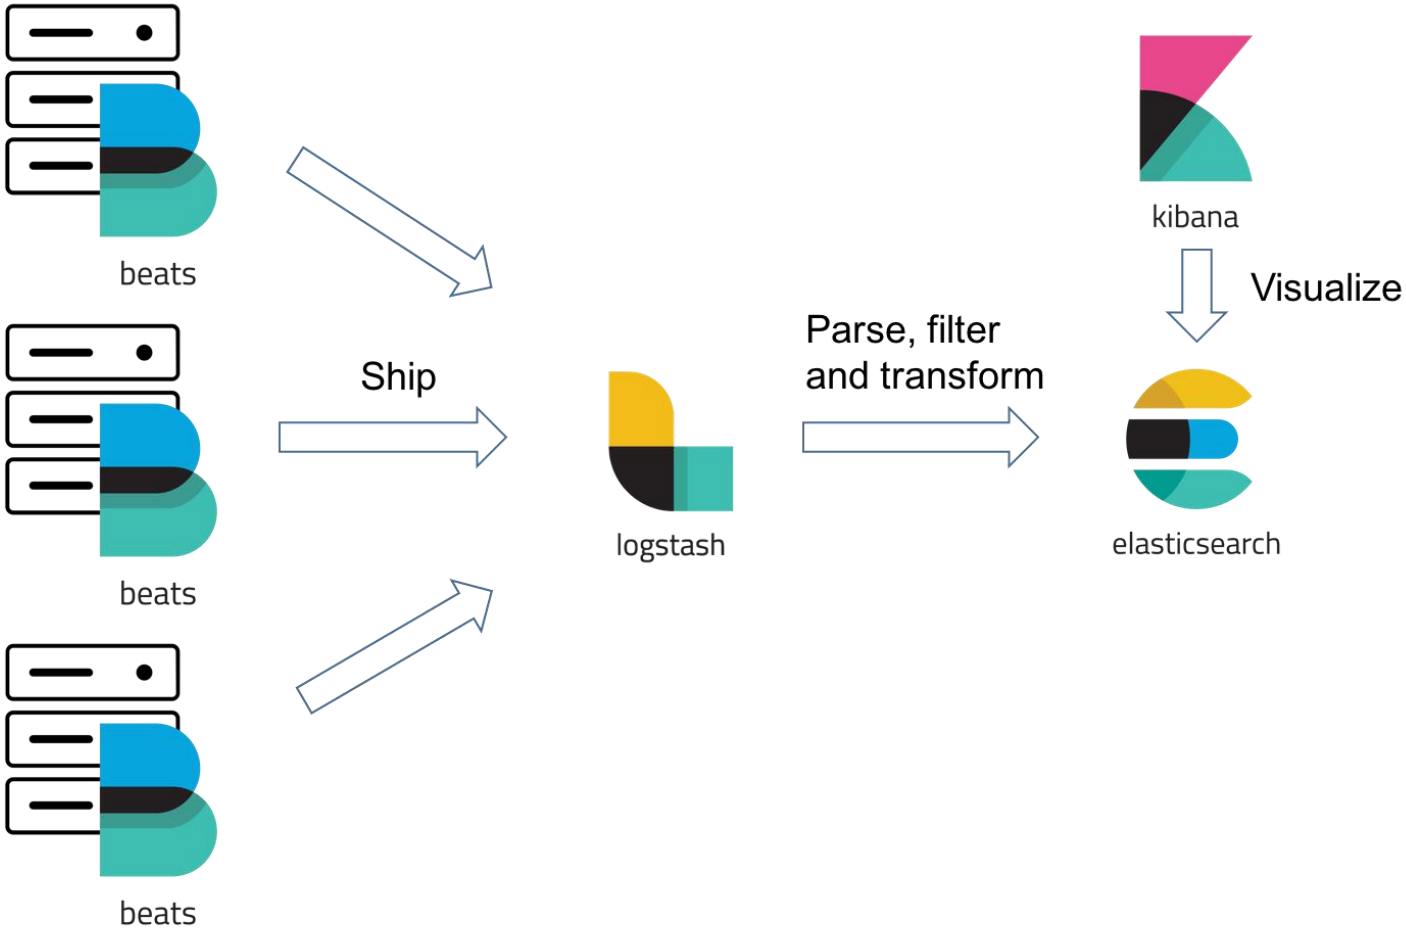
beats
beats
beats
Ship
logstash
Parse, filter
and transform
kibana
Visualize
elasticsearch

Approach

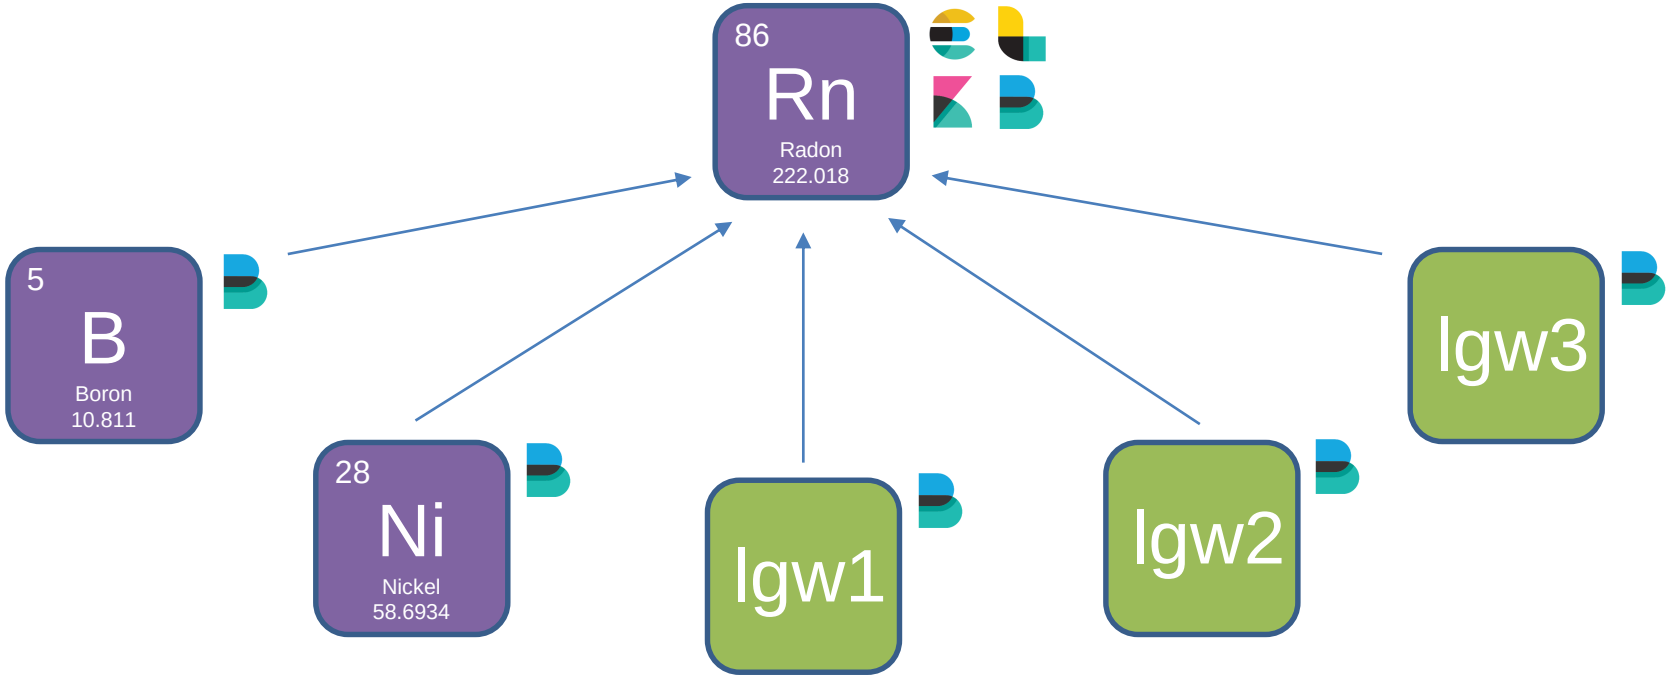
86
Rn
Radon
222.018
5
B
Boron
10.811
28
Ni
Nickel
58.6934
lgw1
lgw2
lgw3

Number of Documents per Node

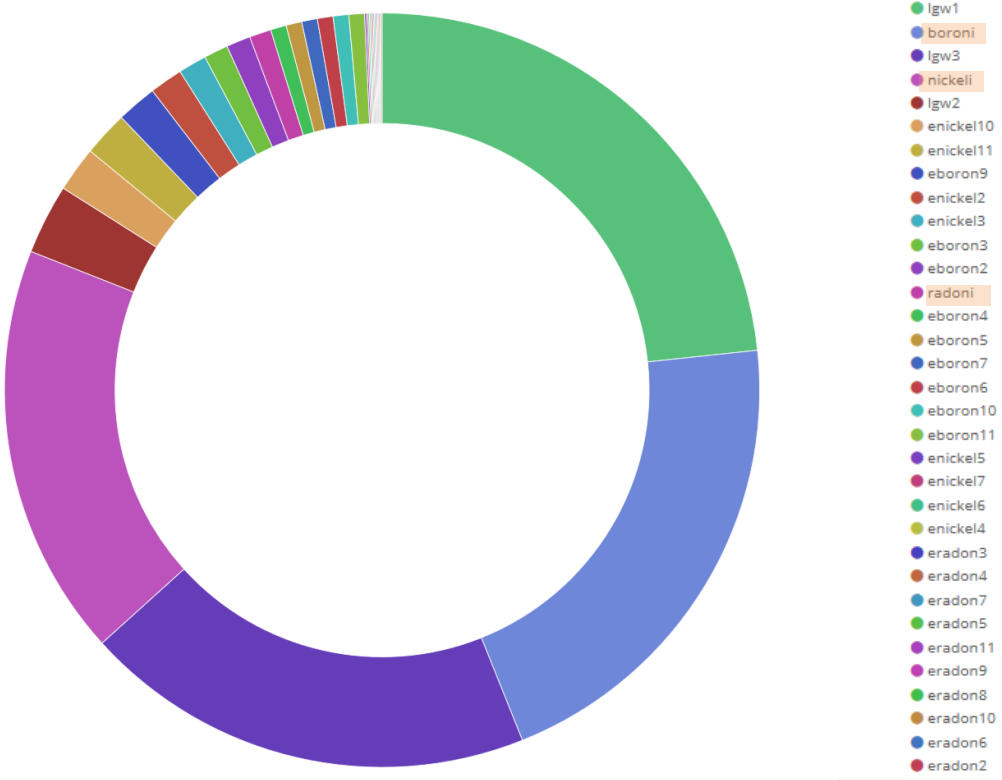

Failed Login Attempts

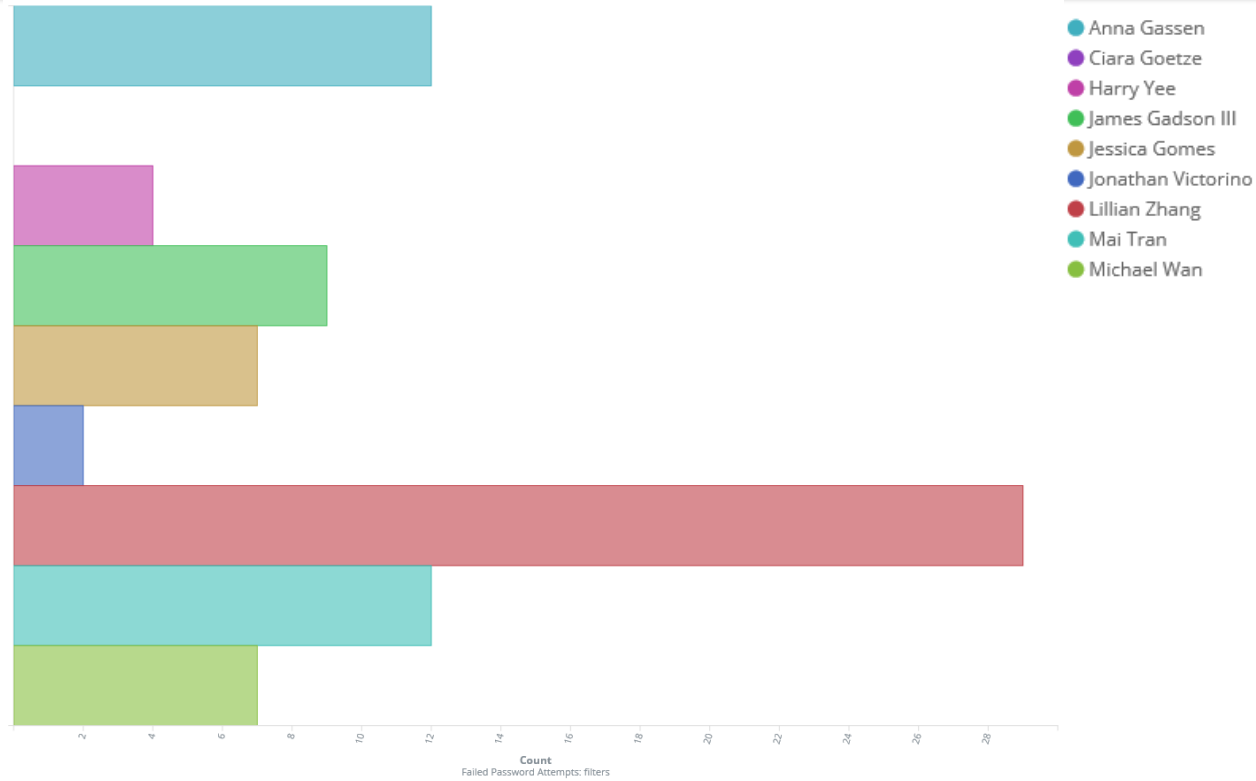

Root vs Non-Root Logins

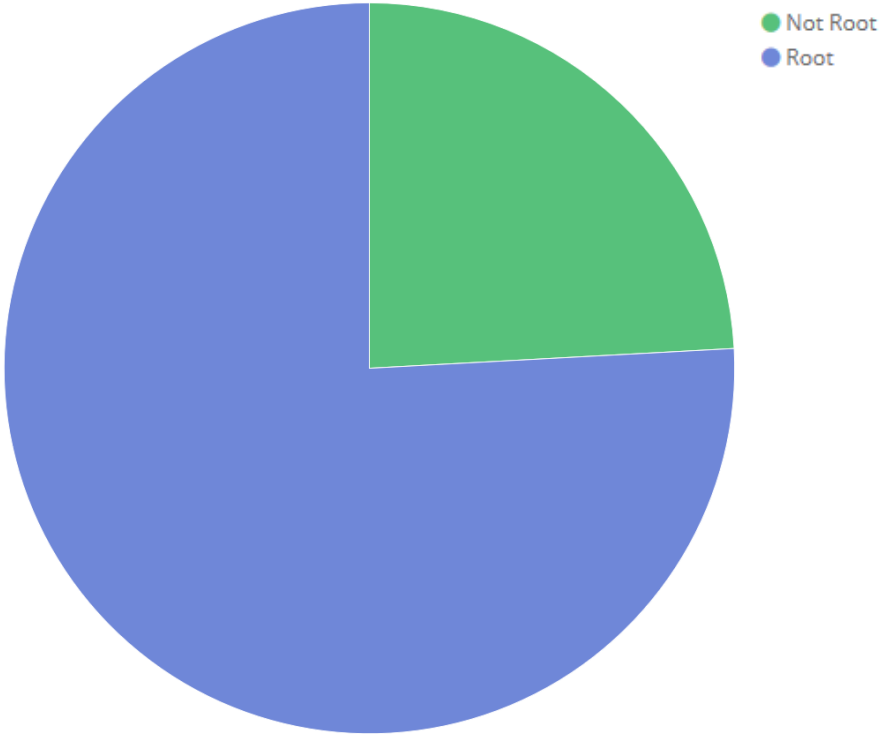

Number of Documents per Day

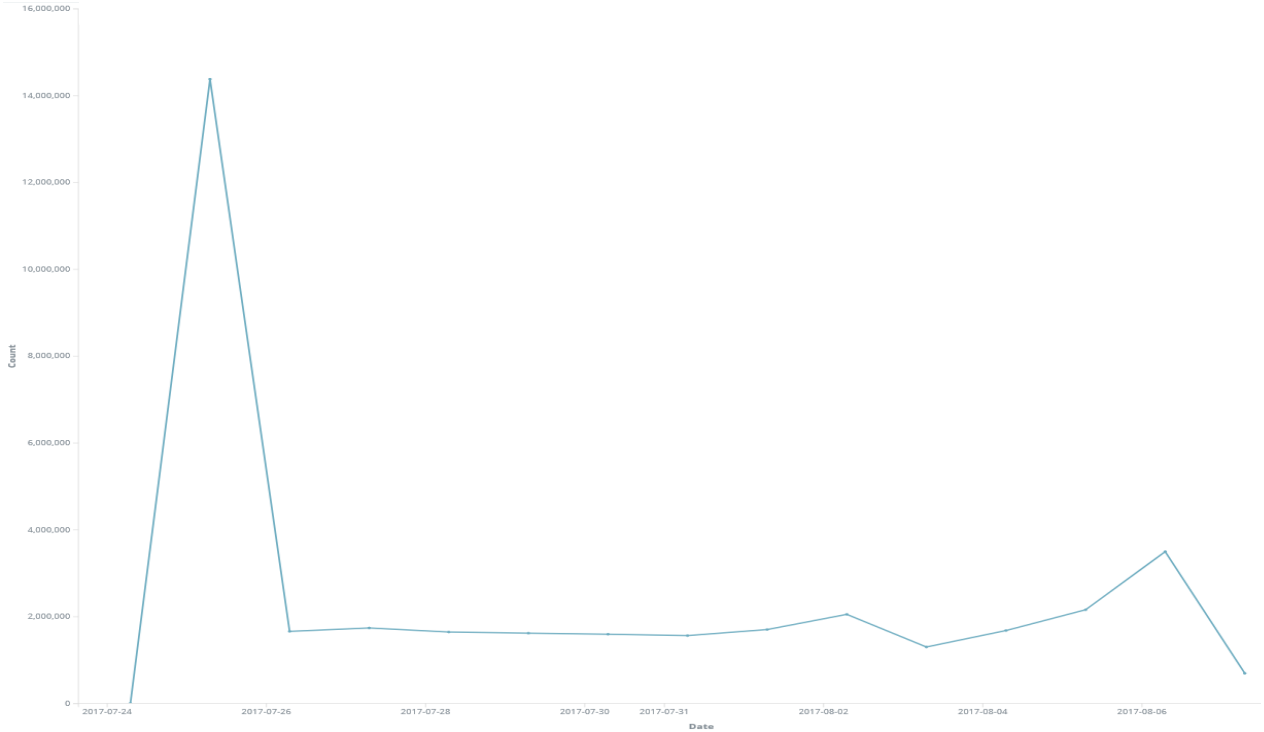

Martian Source Warnings

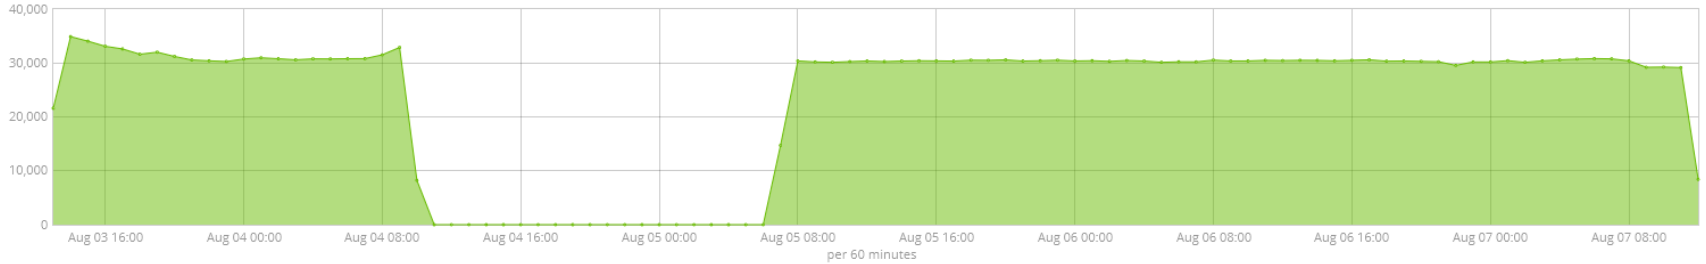

August 7th 2017, 13:22:09.461 @timestamp: August 7th 2017, 13:22:09.461 offset: 35,971,095 @version: 1 beat.hostname: lgw1 beat.name: lgw1 beat.version: 5.5.1 input_type: log hos t: lgw1 source: /var/log/kern.log message: Aug 7 13:22:23 elgw3 kernel: IPv4: martian source 192.168.36.2 from 192.168.36.254, on dev mgmt type: log tag s: beats_input_codec_plain_applied _id: AV2-XKtGgNBZ6calhzUk _type: log _index: filebeat-2017.08.219 _score: -

Future work

- Research Logstash pipeline configuration options
- Utilize Beats and X-Pack
- Perform more complex Elasticsearch queries
- Configuring Elastic Stack to be useful to future Academy interns

Acknowledgements

- David Fox
- Geoff Cleary
- Pam Hamilton
- Bryan Dixon
- Richard Randall


Lawrence Livermore
National Laboratory