Instrumenting and Monitoring TIER Components--First Steps in a Long Journey

Initial Investigations: Analysis of Grouper Logs using the Free and Open Source ELK stack (Elasticsearch, Logstash and Kibana)

Current work: Testing the ELK stack for Grouper log aggregation, exploration and visualization; The Grouper Demo Server is serving as the data source.

The ELK stack brings together

- Logstash for aggregation of monitoring data: Logstash can consume data from multiple sources and transport modes:

syslog	slf4j
log4j	SNMP TRAPS
JMX	JMS
Amazon SNS	Graphite
shell commands	HTTP
RSS	STOMP
IMAP	...and many more

- Elasticsearch for searching, persistence and analysis; Uses https://lucene.apache.org/ as its search engine.
- Kibana for data exploration, analysis and visualization

This whole field is in an explosive phase of growth, and it is well on its way to becoming a first-order discipline of its own. Since TIER architecture is premised on distributed components that are loosely coupled, the ELK stack is an extremely useful tool for addressing TIER monitoring needs.

One of the deep challenges for TIER will be determining which specific sorts of data and data analysis will give us the greatest degree of visibility into the physiological processes and state of health of the distributed TIER infrastructure.