SAND2015-10324PE

Exceptional service in the national interest

Sandia National Laboratories

Firewheel – A Platform for Cyber Analysis

Kasimir Gabert, Adam Vail, Ian Burns, Mack McDonald, Steven Elliott, John Montoya, Jenna Kallaher, Todd Jones, Tan Thai

U.S. DEPARTMENT OF ENERGY | NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2013-XXXXP

Background

- DoD research organization had a need for a very large-scale virtual testbed of 100K+ virtual machines to emulate an internet-like environment and perform large-scale analytics
- Original requirements
  - Scale up to 100K+ end points running Windows & Linux
  - Fast set up and boot time
  - Perform large-scale data analytics
- No national testbed met these requirements
- Initial Firewheel technology was developed for a specific DoD use and under highly compressed schedule
- Redesign and rewrite of Firewheel for stability and validation has occurred under internal research funding since June 2013
- Firewheel has since been used on numerous projects for a variety of customers

Example Questions Firewheel can Answer

Sandia National Laboratories

- How well does a network protocol perform under various adverse network conditions and different configurations?
- What is the impact of a large scale, persistent DNS amplified attack on an enterprise network? How effective are various mitigation strategies?
- What is the potential impact to a network of interest when its own AS de-peers with other ASes?
- Can we detect hidden services in Tor networks?
- How does a botnet or malware behave differently under different network topologies, configurations?
- Can we characterize the performance of a suite of cyber tools (e.g. IDS, firewall, network scanners)?

Emulytics Process

Question
to Answer

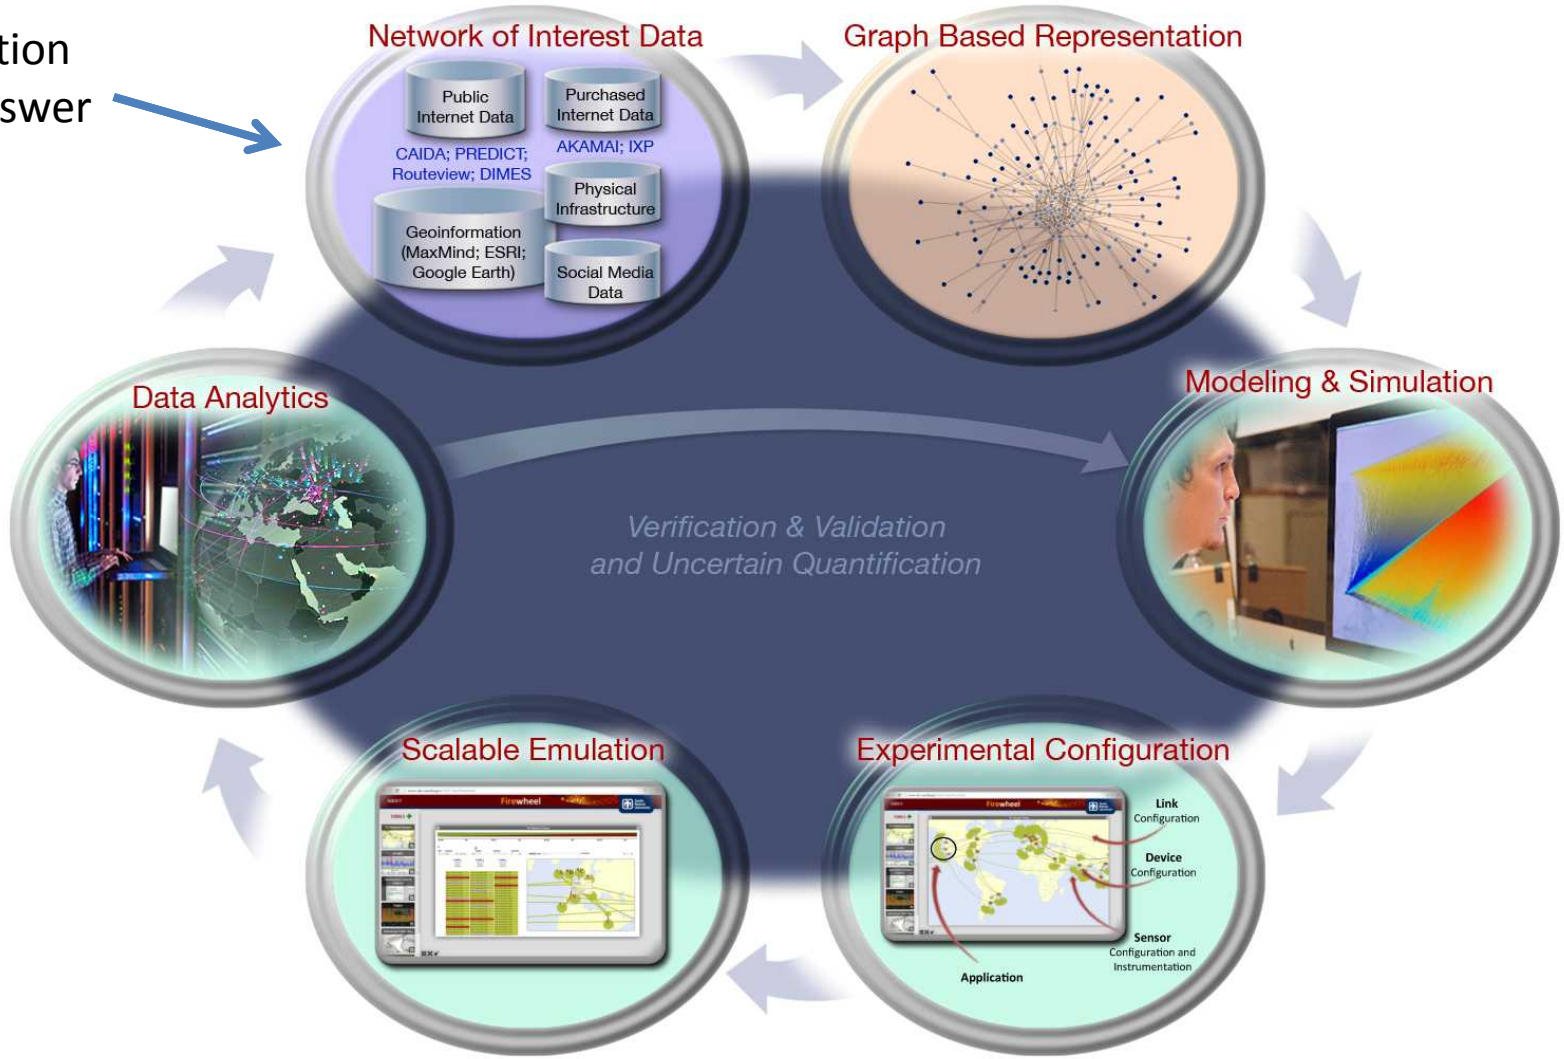

Example Topology: Energy Sciences Network

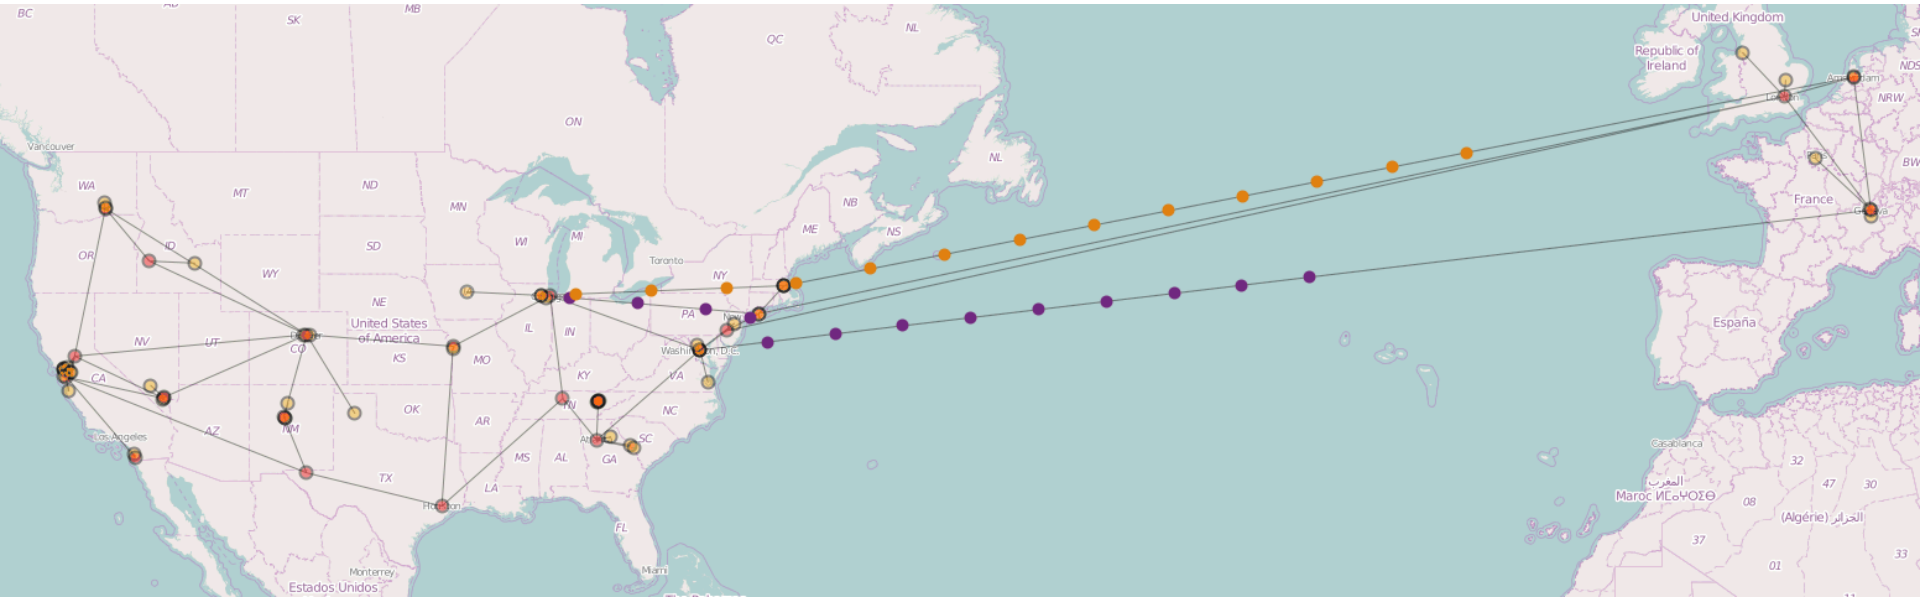

Example Topology: Global DDoS Attack

Sandia National Laboratories

Example Topology: Small Notional Topology

ACME Corporation Network

4 Physical buildings: 3, 2, 7, 10
Building 7: Internet access / IT staff
Building 10: Data center
Building 2: Staff, with critical staff in basement
Building 3: Executive team

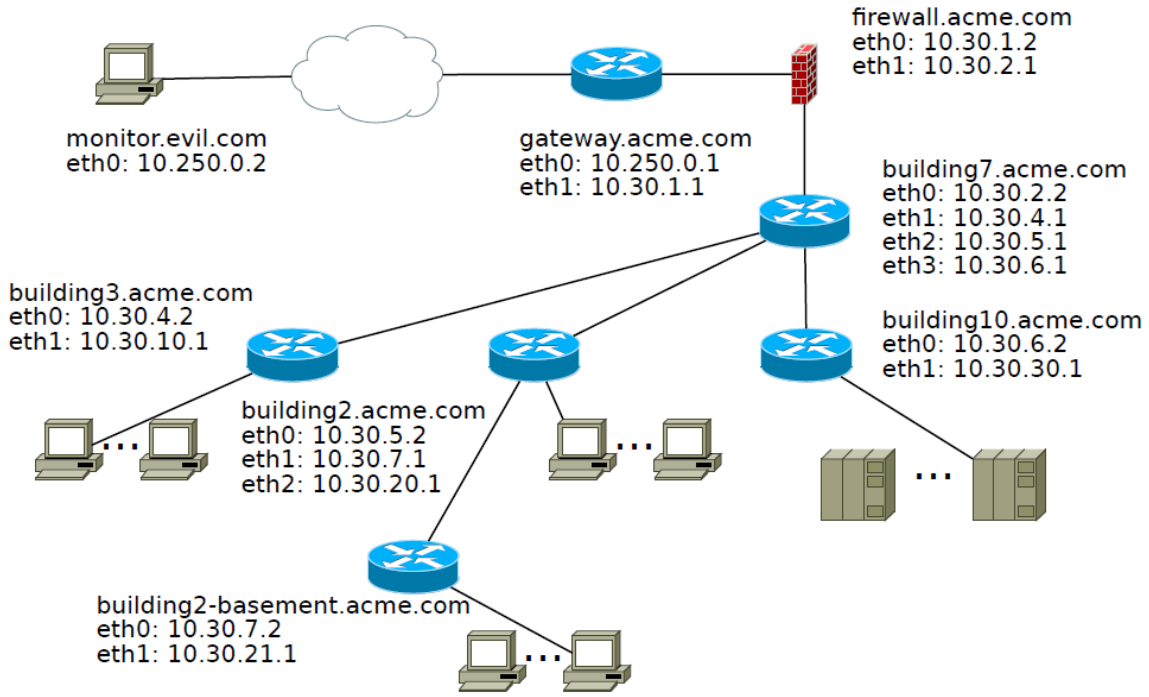

Example Topology: Large AS and its Exit Routers

Sandia National Laboratories

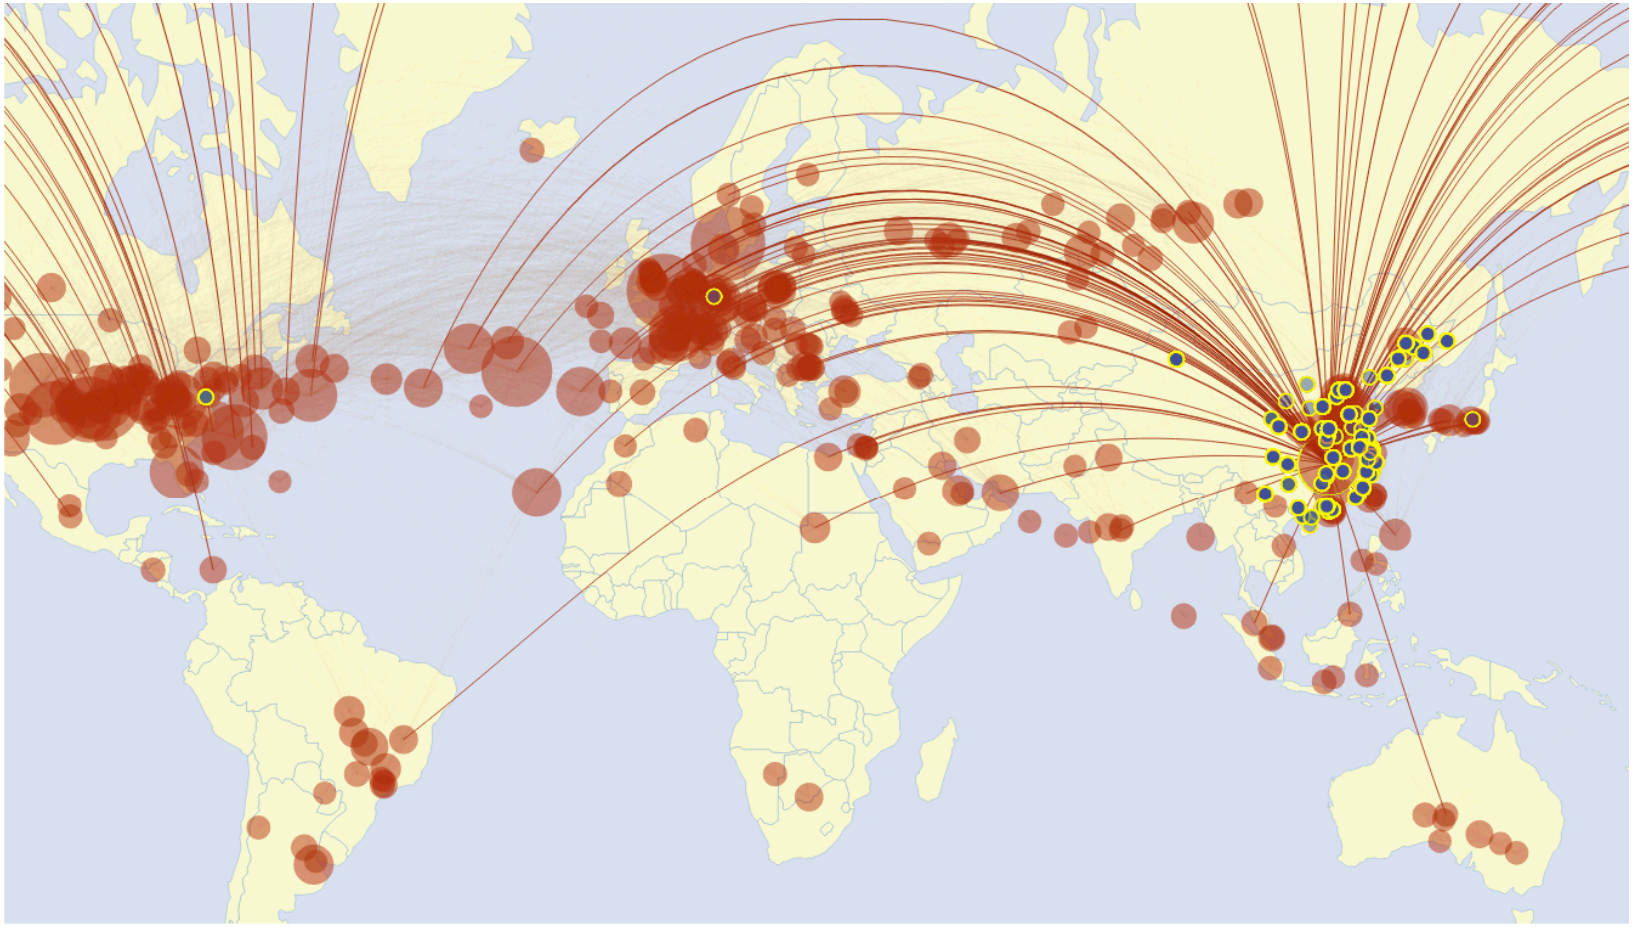

97.249.232.221.broad.wh.hb.dynamic.163data.com.cn (AS4134)

Example Topology: Real-world Tor Network

Sandia National Laboratories

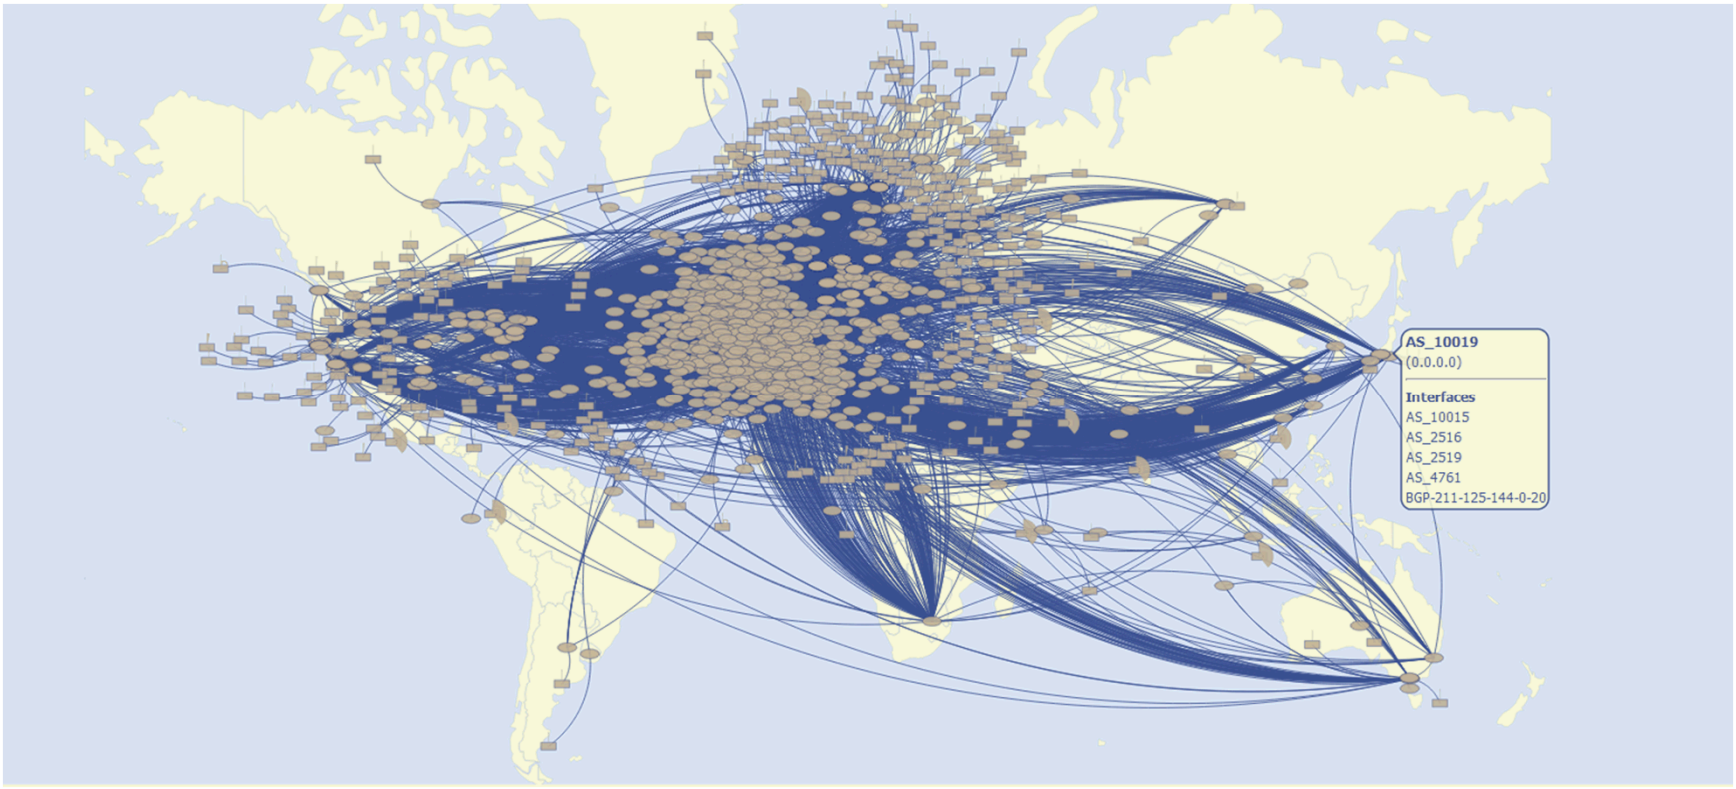

Firewheel Overview

- Large-scale emulation with high fidelity and lower operational cost
- Verification & validation of testbed operational integrity
  - Extensive physical and virtual system monitor and logging
  - Ongoing LDRD work on scientific approaches to this problem
- Support different types of cyber infrastructure
  - Live, Virtual, Constructive (LVC) elements
  - Computer networks, SCADA systems
  - Future development to emulate mobile devices and cellular infrastructure

High-Level Firewheel Components

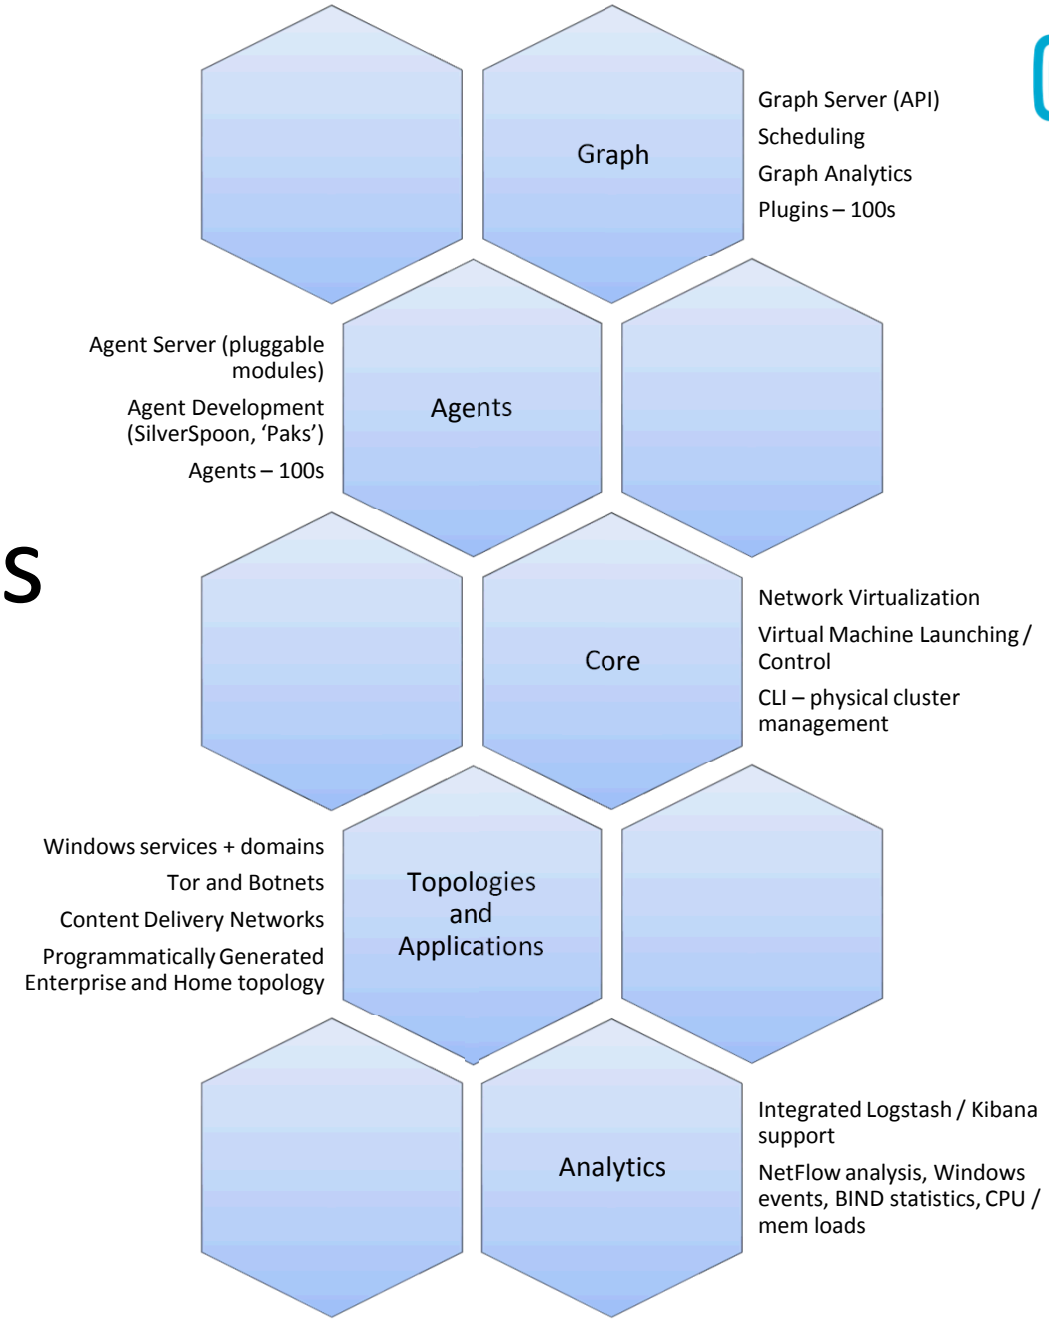

Architecture Efficiency

- Very high density of VMs per core, fast set up and near linear boot time
  - 750 Linux VMs per 16 physical core on Dell C6220
  - 15 minutes to boot 72K Ubuntu 12.04 VMs with full network convergence

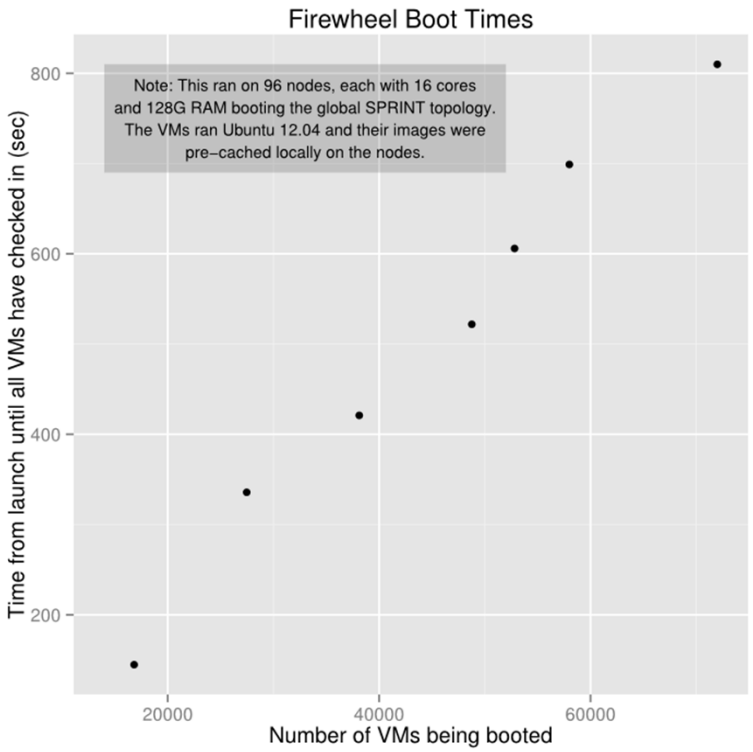

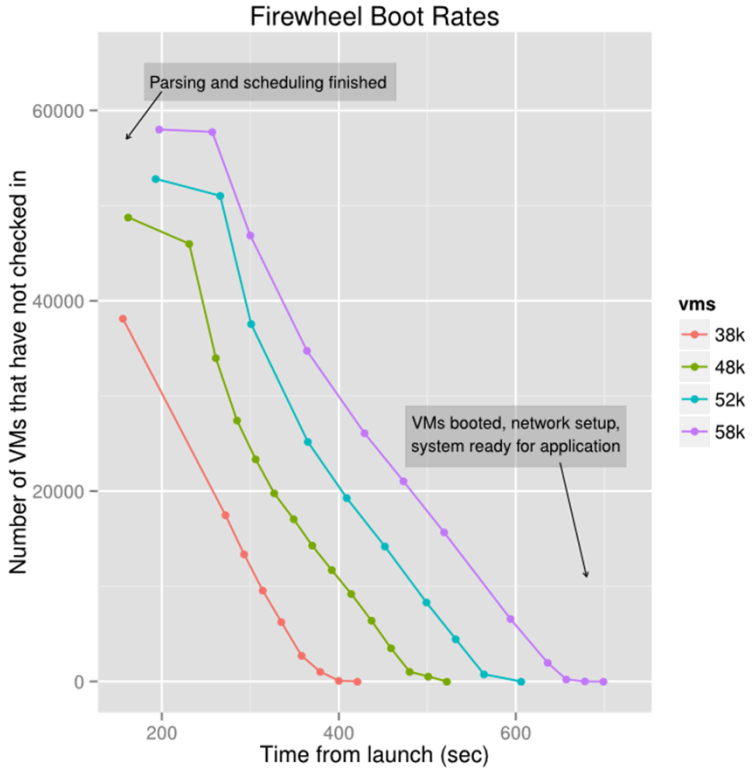

Sandia National Laboratories

Unique Characteristics

- Emulating Internet-like network conditions
  - Faithful routing protocols (BGP, OSPF, upcoming iBGP, MPLS)
  - Configurable communication link quality including jitter, latency, fault injection, packet drop rate, etc
  - Network QoS from real data sources, e.g. Sprint, CAIDA/ARK/traceroute
- Data collection, processing and analytics
  - Real-time state-of-health of both physical and virtual devices
  - System (hosts & VMs) logs, flow records, routing tables, network traffic and data exchanged among test nodes
  - Data analytics with Logstash, ElasticSearch and Kibana
- Lower operational cost
  - High density means lower hardware and operational cost
  - Partition hardware cluster for many concurrent experiments at various scales
  - All open source and no software license fees

Example Physical Testbed

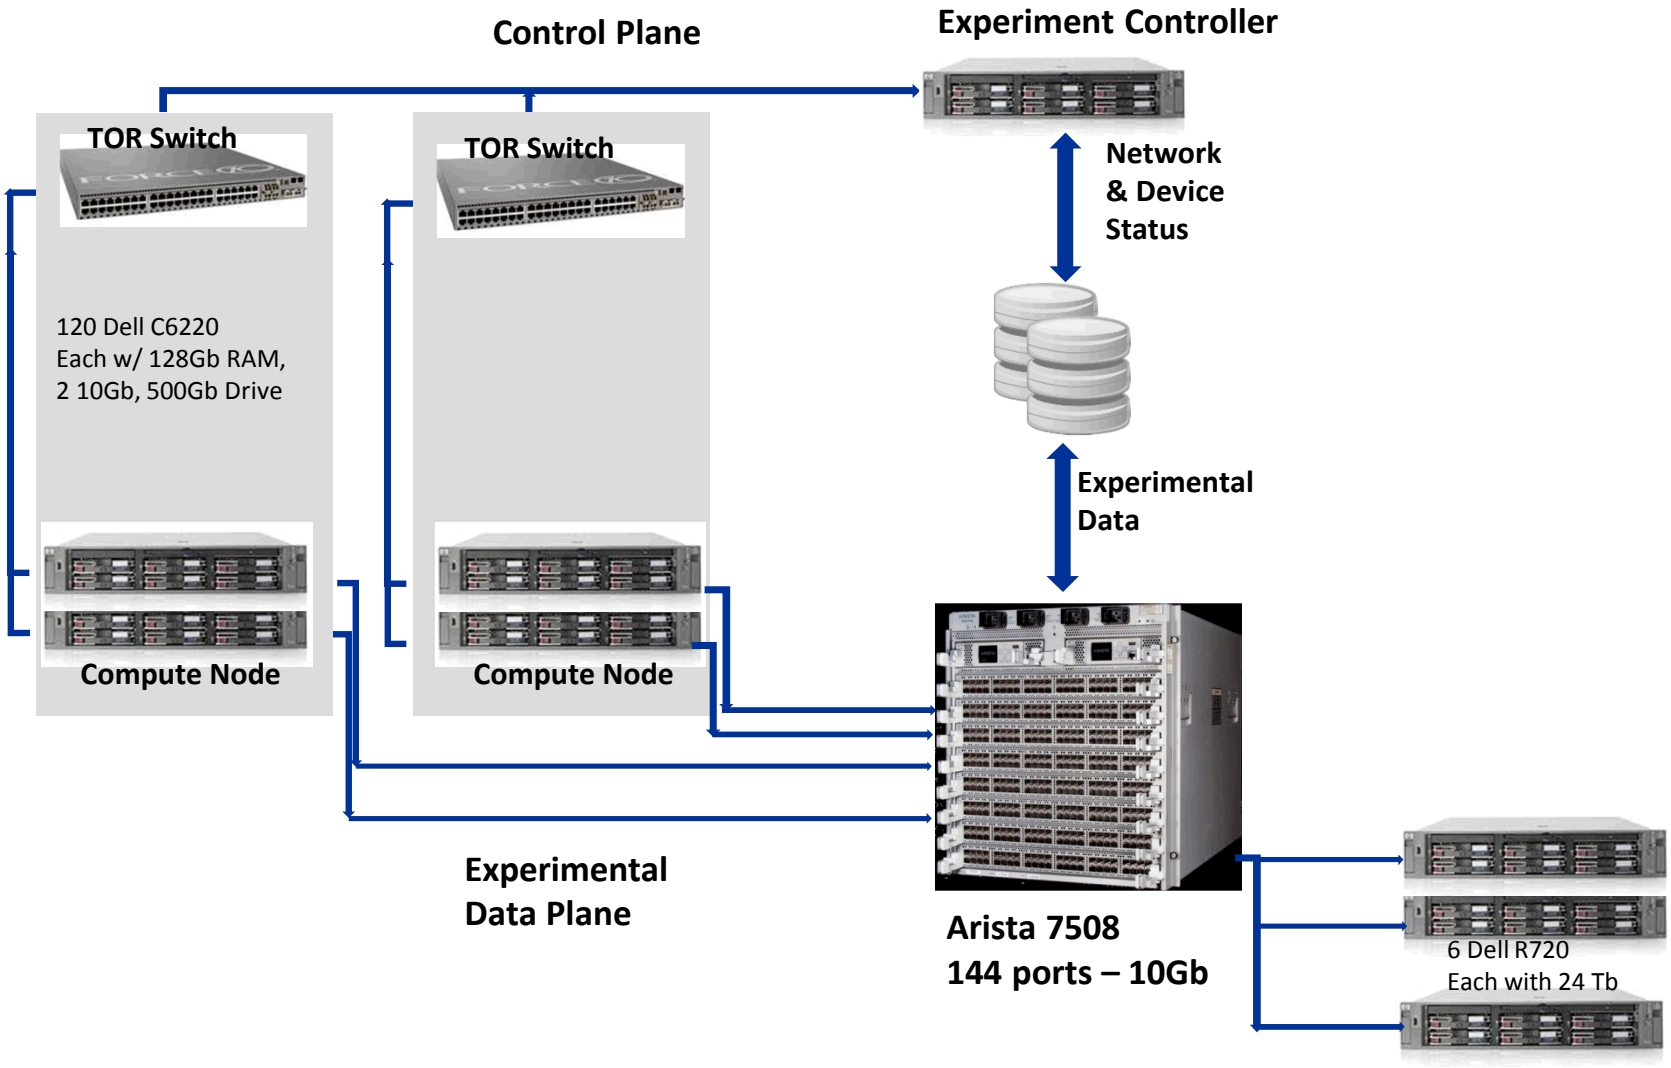

Sandia National Laboratories

Computing Facility

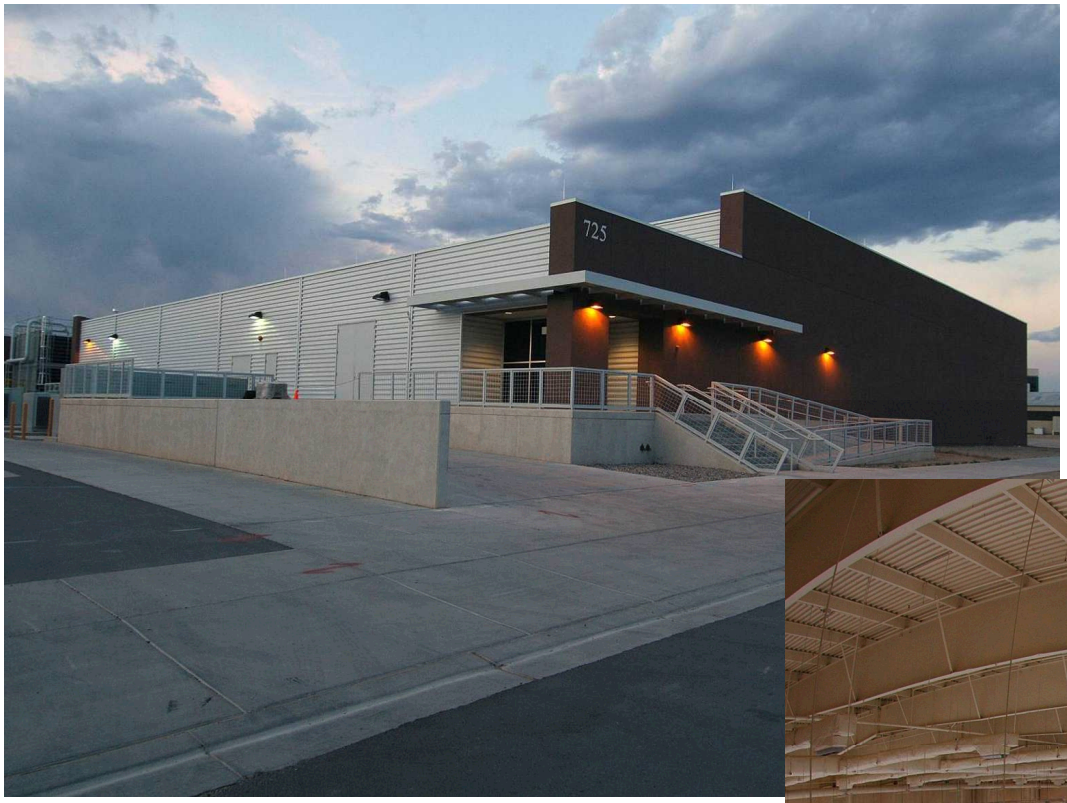

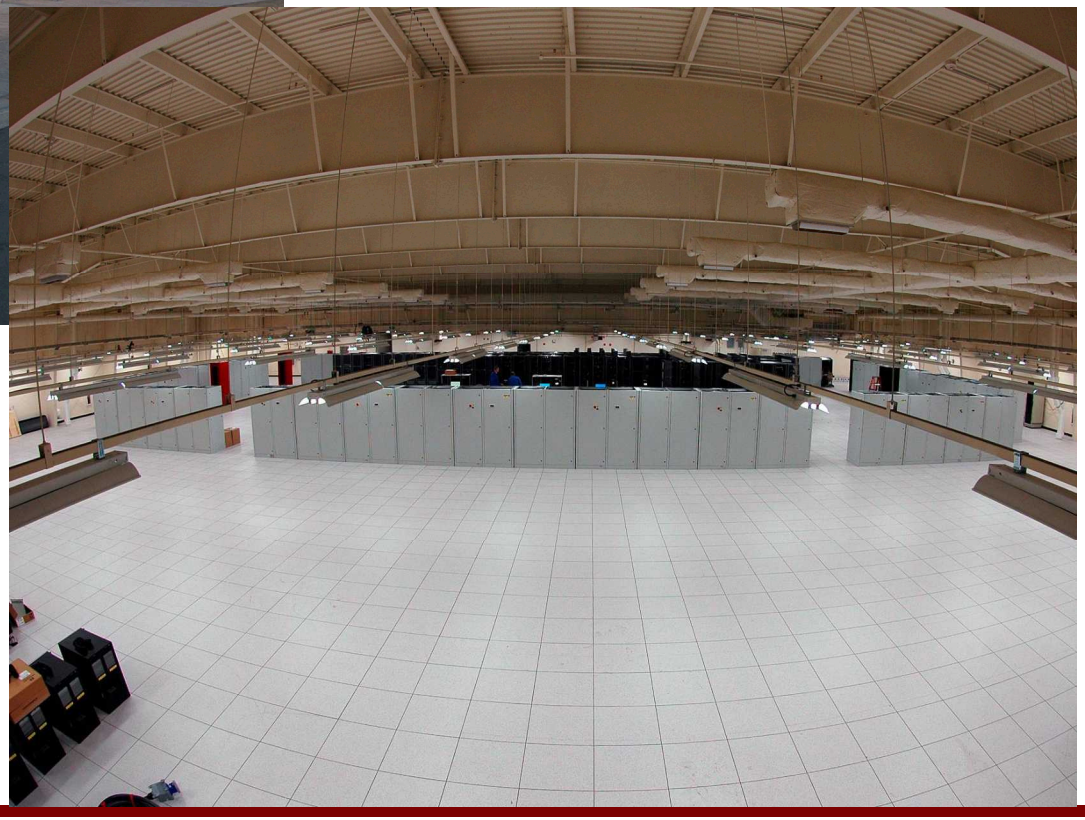

Architecture

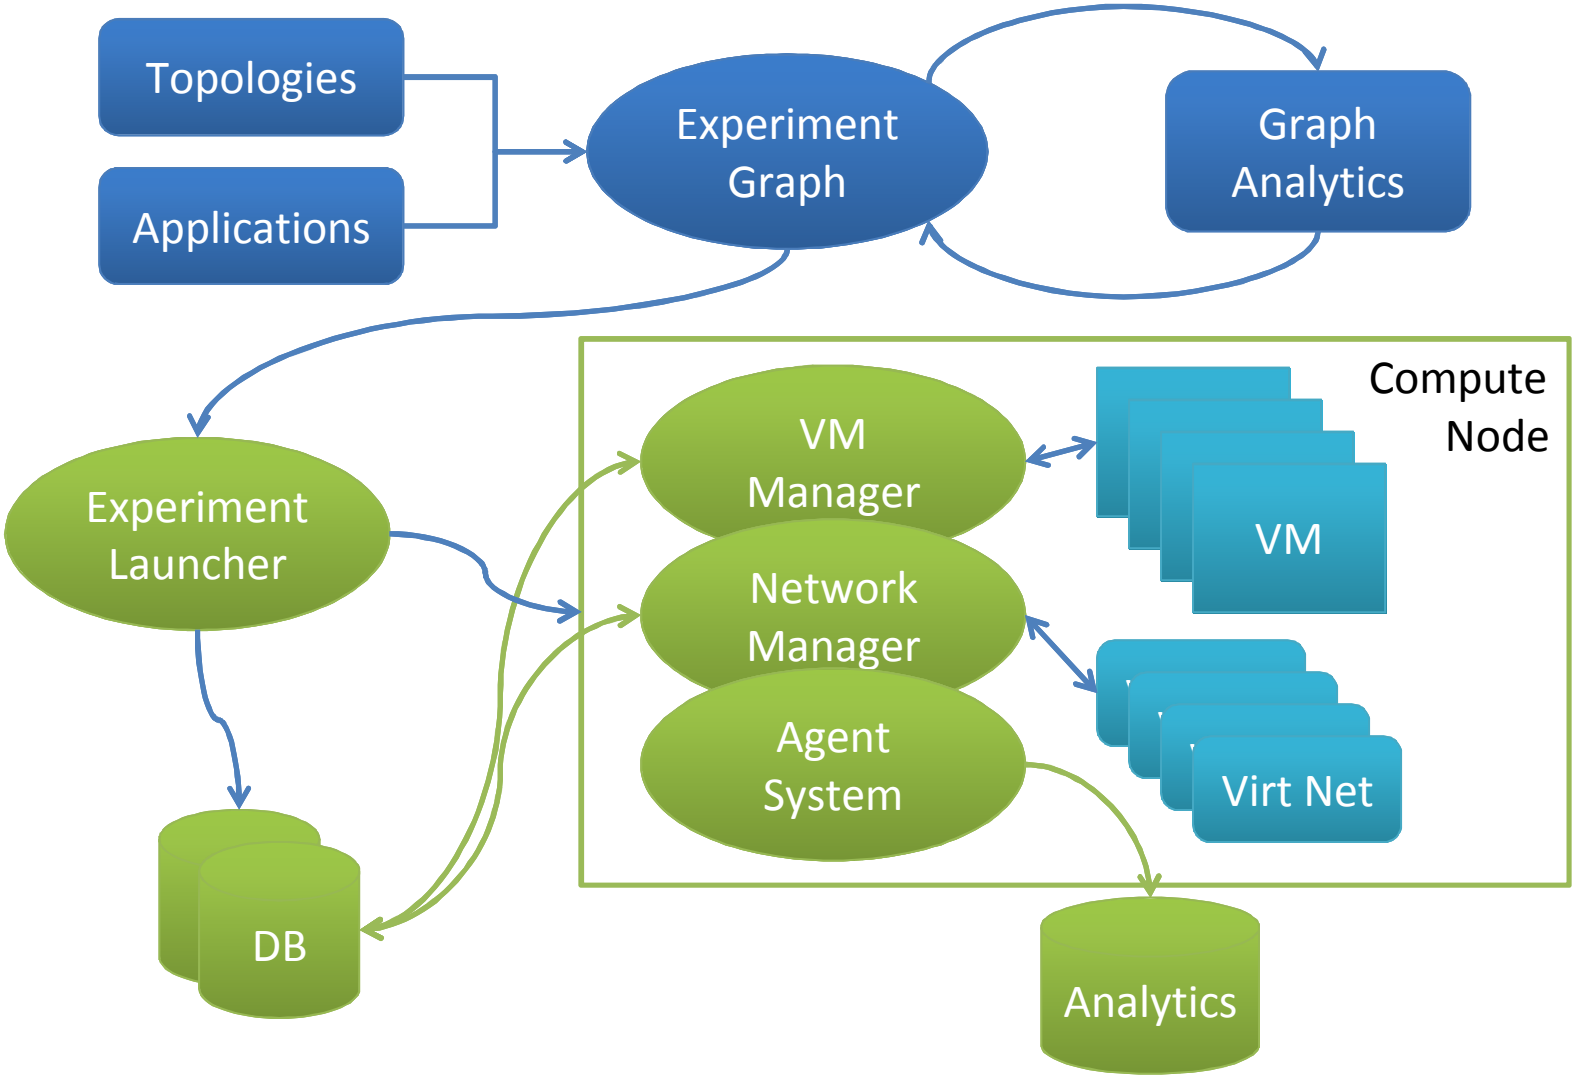

Experiment Graph

- Parses network topology and represents it as graphs
  - Natural network topology representation
  - Dynamic updates of network topology changes during experiment
  - Highly scalable to current and future needs (10M+ nodes and links)

- Modular plug-in model
  - Simple to add support for new network protocols & components
  - Offers dynamic updates via graph streaming
  - Hardware-in-the-loop

Networking Support

- Support for complex virtual networks
  - Logical separation of network traffic via VLAN or VXLAN on physical switches
  - Support for Layer 3 (IP addresses, routing protocols, etc.) in the experiment
  - Support for Layer 2 (VLANs, SDN, etc.) in the experiment
  - Extensible design for virtual networks
- Control plane
  - For state-of-health, VM status, instrumentation data, etc.
  - All VMs have interfaces to control plane
  - Routed via OSPF for scalability (i.e. avoiding ARP table explosion)

Virtual Machine Support

- Uses KVM and controls VMs through QMP APIs
- Leverages Kernel Samepage Merging (KSM) for high density
  - 750 VMs per physical host (full Ubuntu with 256 MB of ram)
  - Better performance with custom Parallel KSM (pKSM) kernel modifications

- Controls virtual machines through a scalable and flexible Agent system
  - Allows virtual machines to execute binaries at a given point in the experiment
  - Allows for dependencies for installs, including rebooting virtual machines, after the experiment has been started
  - Promotes maintainable and reusable experiment configuration

Visualization and Analytics

- Firewheel dashboard for visualization
  - Experiment configuration
  - Network topology display with geographic plotting by IP address
  - State-of-health of both physical and virtual hosts
- Data collection & analytics
  - Logstash + ElasticSearch (Apache Lucene) + Kibana
- Infrastructure monitoring
  - Ganglia