SRI International

USC Viterbi School of Engineering
Information Sciences Institute

Cybersecurity Experimentation of the Future (CEF): Catalyzing a New Generation of Experimental Cybersecurity Research

David Balenson and Laura Tinnel, SRI International
Terry Benzel, USC Information Sciences Institute (ISI)

August 11, 2017

Funded by NSF under Grants No. #ACI-1346277 and #ACI-1346285

Research Infrastructure for Cybersecurity Research

- Cybersecurity R&D is still a relatively young field
- It involves intrinsically hard challenges
  - Inherent focus on worst case behaviors and rare events
  - In the context of multi-party and adversarial/competitive scenarios
- Research infrastructure is crucial
  - Allow new hypotheses to be tested, stressed, observed, reformulated, and ultimately proven before making their way into operational systems
- Ever increasing cyber threat landscape demands new forms of R&D and new revolutionary approaches to experimentation and test
- Clearly a need for future research infrastructure that can play a transformative role for future cybersecurity research

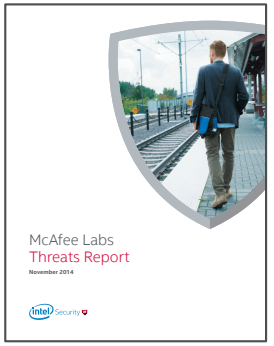

CEF Study

Community-based effort to study current and expected cybersecurity experimentation infrastructure, and to produce a strategic plan and enabling roadmap intended to catalyze a new generation of experimental cybersecurity research

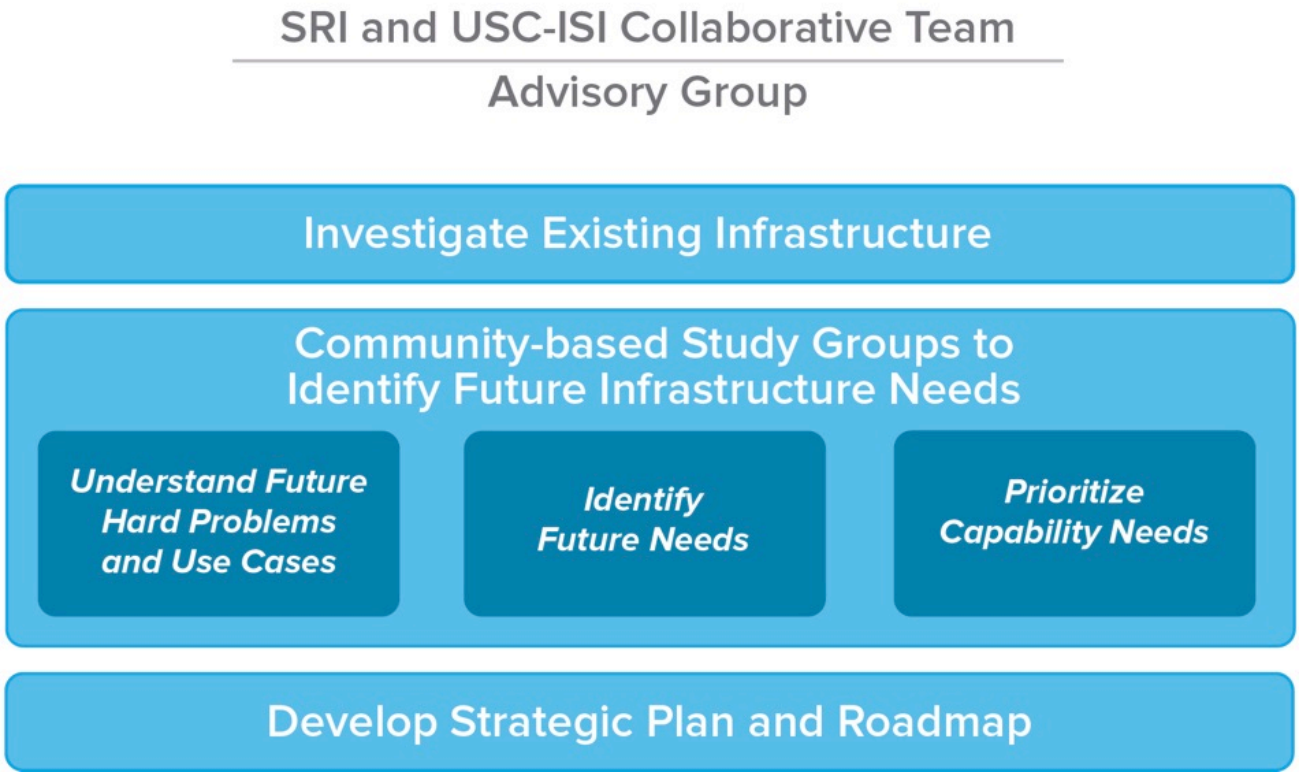

CEF Study Results

CEF Final Report published in July 2015

- Vision for cybersecurity experimentation infrastructure
- Core capabilities and research needed to achieve them
- Survey of existing infrastructure and high-level gap analysis
- Top five recommendations
- Overarching findings

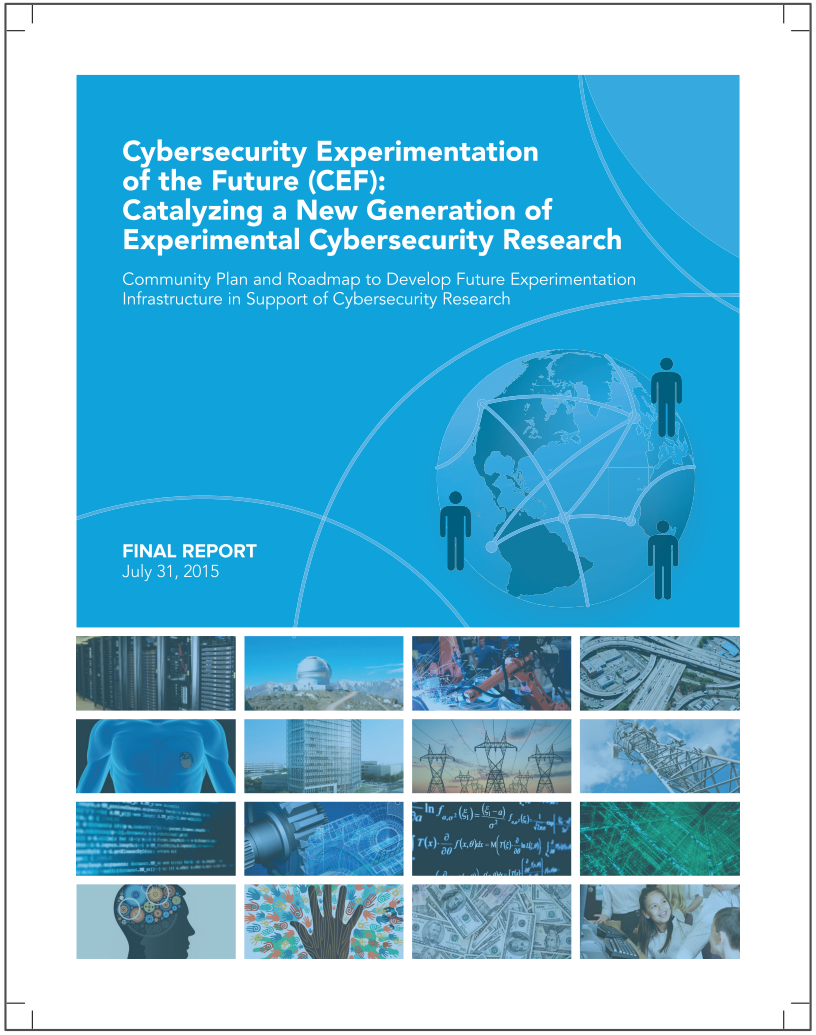

http://www.cyberexperimentation.org/

The Need for Transformational Progress

Transformational progress in three distinct, yet synergistic areas is required to achieve the desired objectives:

1) Fundamental and broad intellectual advances in the field of experimental methodologies and techniques
   - With particular focus on complex systems and human-technical interactions
2) New approaches to rapid and effective sharing of data and knowledge and information synthesis
   - That accelerate multi-discipline and cross-organizational knowledge generation and community building
3) Advanced experimental infrastructure capabilities and accessibility

A Science of Cybersecurity Experimentation!

Science of Cybersecurity Experimentation

- New direction for the field of experimental cybersecurity R&D
- R&D must be grounded in scientific methods and tools to fully realize the impact of experimentation
- Different than and complementary with the science of cybersecurity

Source: https://www.nsa.gov/research/tnw/tnw192/article4.shtml

- New approaches to sharing all aspects of the experimental science – data, designs, experiments, and research infrastructure
- Cultural and social shifts in the way researchers approach experimentation and experimental facilities
- New, advanced experimentation platforms that can evolve and are sustainable as the science and the community mature

Roadmap for a New Generation of Experimental Cybersecurity Research

- The roadmap presents requirements, objectives and goals for 30 key capabilities organized into 8 core areas over 3, 5, and 10 year phases
  - Some phases build upon each other and others require new fundamental research over a long time period

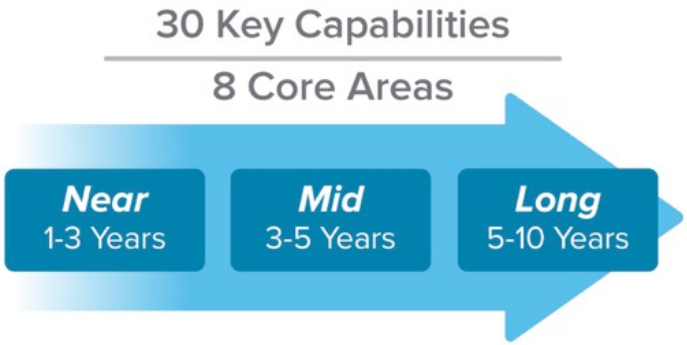

- Key capabilities consider:
  - Current experimental cybersecurity research and its supporting infrastructure
  - Other types of research facilities
  - Existing cyber-domain “T&E” capabilities (primarily DoD)
- The roadmap presumes advances in key computer science disciplines
  - Ontologies, meta-data, libraries, and corresponding resource discovery

Core Capability Areas

8 core areas layered from the outside “application” layer down to the base system and a corresponding set of meta-properties

4.1 Domains of applicability

4.2 Modeling the real world for scientifically sound experiments

4.3 Frameworks and building blocks for extensibility

4.4 Experiment design and instantiation

4.5 Interconnected research infrastructure

4.6 Experiment execution and management

4.7 Instrumentation and experiment analysis

4.8 Meta-properties

Top 5 Recommendations

Domains of Applicability – Multidisciplinary Experimentation: Focus on multidisciplinary experimentation that includes computer science, engineering, mathematics, modeling, human behavior, sociology, economics, and education

Modeling the Real World for Scientifically Sound Experiments – Human Activity: Accurately represent fully reactionary complex human and group activity in experiments, including live and synthetic humans

Frameworks and Building Blocks for Extensibility – Open Interfaces: Create open standards and interfaces, for both experimental infrastructure facilities and for experiments themselves

Experiment Design and Instantiation – Reusable Designs for Science-based Hypothesis Testing: Experiment designs and design patterns for designing meaningful experiments that reflect the real world

Meta-properties – Usability and Cultural Changes: Cybersecurity research infrastructure must be usable by a wide range of researchers and experts across many different domains of research, and researchers must make a concerted effort to take advantage of community based resources

A Definition of "Cybersecurity Experimentation Infrastructure"

- General purpose ranges and testbeds (physical and/or virtual)
- Specialized ranges and testbeds (physical and/or virtual)
- Software tools that supports one or more parts of the experiment life cycle, including, but not limited to:
  - Experiment design
  - Testbed provisioning software
  - Experiment control software
  - Testbed validation
  - Human and system activity emulators
  - Instrumentation – systems and humans
  - Data analysis
  - Testbed health and situational awareness
  - Experiment situational awareness
  - Other similarly relevant tools


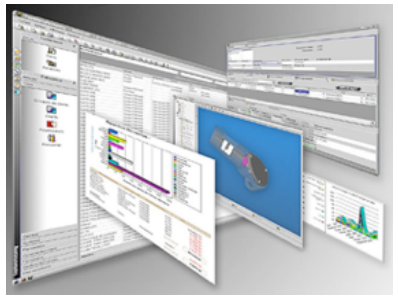
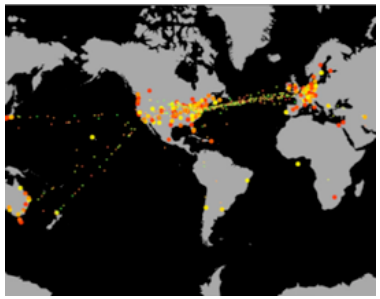


- Specialized hardware tools – simulators, physical apparatus, etc.

Ecosystem of Different Experimental Capabilities Spanning Multiple Domains

- The goal is not to create a single instance of a cyber experimentation testbed or facility
- Over time the roadmap may be realized through an ecosystem of many different instantiations – from small, stand-alone and localized to large distributed experimental capabilities, all spanning multiple domains

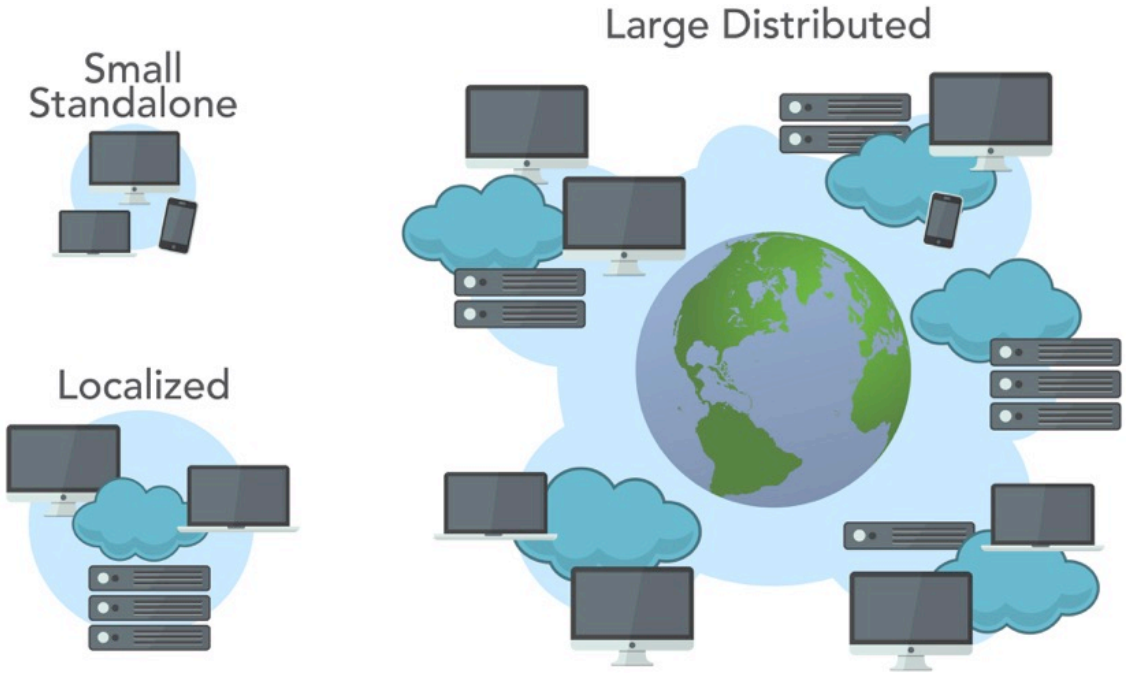

Hybrid Architectures Based on Different Building Blocks

- Cloud technology
- Software defined networking (SDN)
- Knowledge sharing and community environments
  - E.g., Collaboration sites and wiki's
- Integrated Development Environments
  - E.g., Eclipse
- Emulated and simulated environments
  - E.g., Real time digital simulator (RTDS), wireless
- Specialized hardware
  - E.g., FPGA, GPU, Intel Xeon Phi

- No single hardware/software substrate

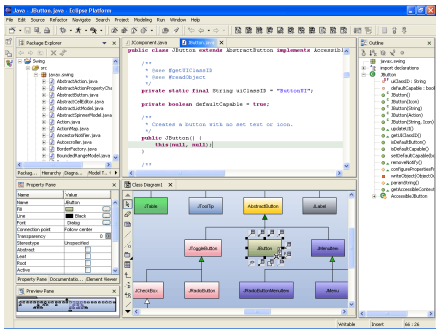

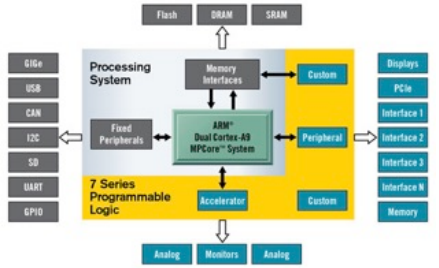

Research Infrastructure is More than Infrastructure

- Research infrastructure **>** infrastructure of machines and tools
  - Scientific methodologies, experimental processes, and education are critical to effective use of the machines and tools
- Research infrastructure requires meta-research into:
  - Design specification (multi-layered languages and visualization)
  - Abstraction methodologies and techniques
  - Semantic analysis and understanding of experimenter intent
  - Formal methods and a rich approach to modeling to satisfy science objectives

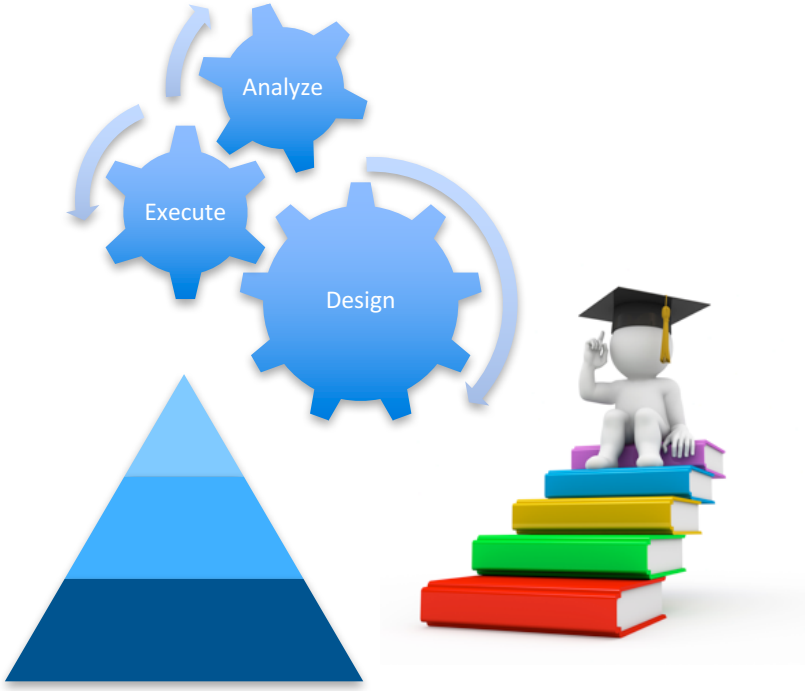

Where is Experimentation Applicable?

- Overarching goal is to increase researcher effectiveness and support the generation and preservation of solid empirical evidence
  - Infrastructure to enable research, not constrain
  - New mechanisms to capture and share knowledge (designs, data and results) to enable peer review and allow researchers to build upon each other
- Experimentation is about learning
  - To perform an evaluation (not formal T&E)
  - To explore a hypothesis
  - To characterize complex behavior
  - To complement a theory
  - To understand a threat
  - To probe and understand a technology

Representative Cybersecurity Hard Problems

- Systems/software
  - Human interactions
  - System of system security metrics
  - Emergent behavior in large scale systems
  - Supply chain and root of trust
  - Societal impacts and regulatory policies
- Networking
  - Anonymity and privacy of data and communication
  - Trust infrastructure
  - Software defined networking (SDN)
  - Political, social, and economic (balance-of-interest) goals in network design
  - Pervasive communications, across organizational and political boundaries

- Cyber-physical systems
  - Embedded devices
  - Autonomous vehicles, smart transportation
  - Electric power, smart grid
  - Medical implants, body sensors, etc.

Experimentation – It's About the Real World

- Experimentation should start with models of the real world
- Modeling and abstraction allow us to capture conceptual models of the real world with varying degrees of fidelity
- Key research areas include:
  - Experiment design specifications
  - Auto-generated model refinement
  - Methodologies and tools to assess representation
  - Understanding the multiple dimensions of realism
  - Taxonomies of realism metrics for realism sufficiency
- Additional methods and tools can help extend these modeling activities to provide increasing forms of real world models

Survey of Existing Infrastructure – Approach

- Focused on cybersecurity research infrastructure in the U.S., but did consider two sets outside the U.S.
- Set of 46 candidates narrowed to a final, representative set of 18 testbeds or tools that are either commonly available or provide significant value
  - Generic testbeds with open source tools
  - Special purpose testbeds with proprietary tools
  - Specific tools that are not part of any particular testbed

Survey of Existing Infrastructure – Summary

- Experimentation infrastructure and capabilities needed to support cybersecurity research either does not exist or is not generally accessible
  - Restricted government use tools could provide additional capabilities
- Many cybersecurity researchers must design and create a full test apparatus from scratch, which consumes time and resources
  - Seldom shared or reused by other researchers
  - Increases the probability of introducing errors
- Shared, vetted community cybersecurity experimentation infrastructure and tools can help reduce the time, resources, and error

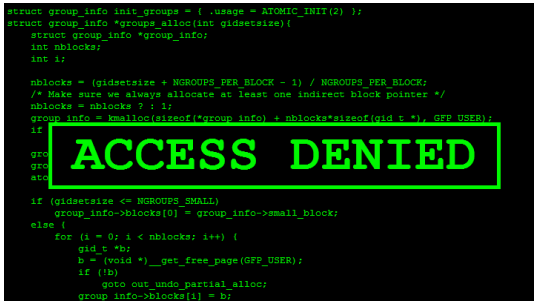

Capabilities Roadmap

For each key capability, to understand the future vision, the current state, and how to achieve the vision, the roadmap asks and answers the following questions:

- *Initial description of the capability*
- *Where do we want to be in the future? What will solutions look like? What will the solutions enable us to do?*
- *Where are we today? What are the current technology and research? What are the important technology gaps and research problems?*
- *What will it take to get us there? How can we divide the problem space? What can be done in the near (1-3 years), mid (3-5 years), and long term (5-10 years)? What resources are needed?*

Core Capability Areas Revisited

4.1 Domains of applicability

4.2 Modeling the real world for scientifically sound experiments

4.3 Frameworks and building blocks for extensibility

4.4 Experiment design and instantiation

4.5 Interconnected research infrastructure

4.6 Experiment execution and management

4.7 Instrumentation and experiment analysis

4.8 Meta-properties

4.1 Domains of Applicability

- Cybersecurity challenges today are very real and wide-ranging with significant implications across many critical sectors of our society
- Cybersecurity experimentation needs to be applicable across multiple domains and communities

- Capabilities:
  - Support for cross domain experimentation (critical infrastructure sectors)
  - Multidisciplinary experimentation (computer science, engineering, math/modeling, human behavior, sociology, economics, education)
  - Portability of experiments, for sharing and use in cross-discipline experiments

4.2 Modeling the Real World for Scientifically Sound Experiments

- To be impactful, cybersecurity research must be based on both sound science and on the real world
- The community needs shared, validated models and tools that help researchers rapidly design meaningful experiments and environments
  - Both real and simulated environments
- Experimentation grounded in the real world is one of the founding principles for the roadmap

- Capabilities:
  - Models of real world environments
  - Experiments that scale
  - Experimentation with systems-of-systems
  - Human activity

4.3 Frameworks and Building Blocks for Extensibility

- Cybersecurity research would most benefit from an ecosystem of many testbeds and tools
  - The goal is not to create a single instance of a cyber experimentation testbed or facility
- Drives need for architectures to provide generic frameworks and specialized instantiations for domain-specific research
  - Infrastructure and experiments

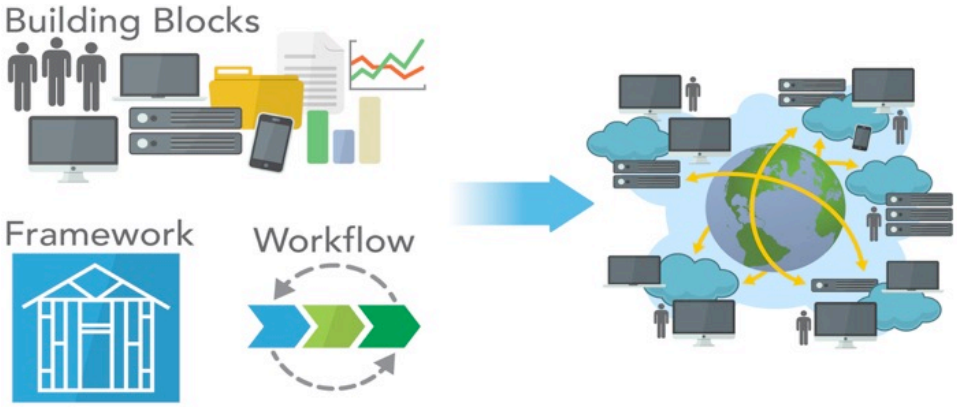

- Capabilities:
  - Workflow and management – comprehensive, human
  - Open, standard interfaces – API for extensibility, plugins write to API
  - Building blocks – libraries
  - Tool integration framework – to glue pieces together

4.4 Experiment Design and Instantiation

- The need for architectures that support multi-domain, multi-discipline cybersecurity research leads to fundamental questions in experiment design and scenario exploration
- Need to move from low-level "realization" mechanisms to higher-level design and representation requirements

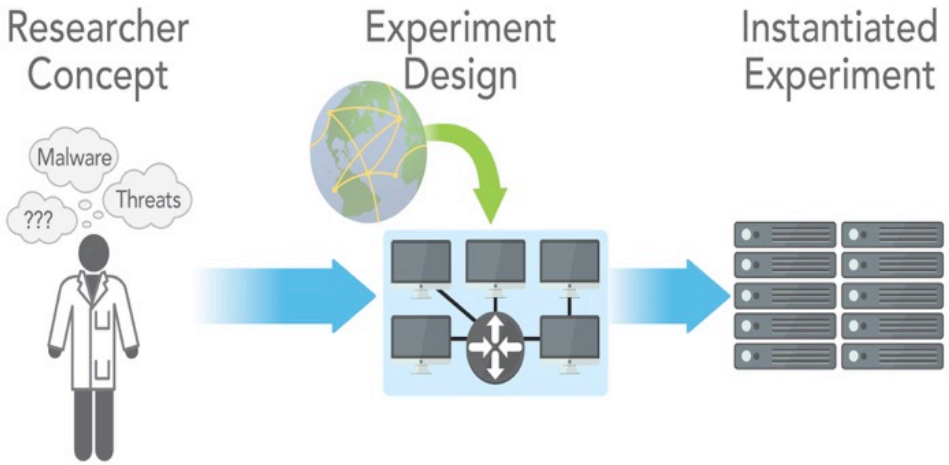

- Capabilities:
  - Design tools, specifications, ontologies, compiler
  - Reusable designs for science-based hypothesis testing
  - Automated discovery of local and distributed resources
  - Dynamic instantiation of domain-specific test apparatus
  - Validation of instantiated test environments and apparatus

4.5 Interconnected Research Infrastructure

- Cybersecurity research requires infrastructure that can support increased scale, complexity, and resource sharing
- Drives the need for flexible and automated capability for interconnection
- Leads to a number of important engineering challenges in connection fabrics

- Capabilities:
  - Automated, transparent federation to interconnect resources
  - Dynamic and on demand, with sharing models
  - Support integrated experiments that include real, emulated (virtual), and simulated elements

4.6 Experiment Execution and Management

- Cybersecurity research requires sound science that builds upon controlled, well-executed experiments
- Drives the need for test environments that allow researchers to intelligently run and control their experiments

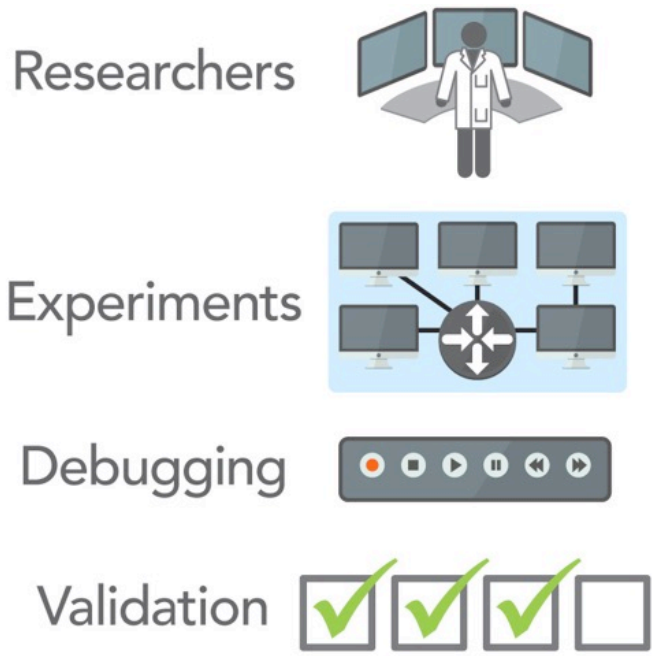

- Capabilities:
  - Experiment orchestration for automation and control
  - Visualize and interact with experiment process
  - Experiment debugging with checkpoint and rollback
  - Experiment execution validation

4.7 Instrumentation and Experiment Analysis

- Cybersecurity research requires the collection, ensured integrity, and analysis of experimental data
- Roughly analogous to inserting debugging code and conducting post analysis of debugging outputs and log files created during execution

- Capabilities:
  - Instrumentation and data collectors
  - Transport and protection mechanisms
  - Data repositories
  - Data analysis

4.8 Meta-Properties

- Research infrastructure needs to be easy to use by a wide range of experimenters and by infrastructure owner/operator
- Community needs mechanisms and processes to provide confidentiality, integrity, and availability of the experiment ecosystem
- Cultural and social changes, along with community building, are needed to facilitate future infrastructure

- Capabilities:
  - Usability (experiments, owner/operator)
  - Confidentiality, integrity, and availability of experiment ecosystem
  - Social and cultural changes

What's Next? Ongoing/Emerging Testbeds

- USC-ISI DETERLab (infrastructure and tools)
- Utah Emulab (Internet, 802.11, SDR)
- DHS CEF Testbed (RFI)
- MIT-LL Cyber Range (and LARIAT, K0ALA, GOSMR, and LO-PHI tools)
- SimSpace Cyber Range
- AFRL Cyber Experimentation Environment (CEE)
- TRMC National Cyber Range (NCR)
- ISU PowerCyber
- PNNL powerNET
- UIUC TCIPG testbed, CRED-C Federated Testbed, and CEER
- FAA Cybersecurity Test Facility (CyTF)
- NHTSA Vehicle Research & Test Center
- ACM Willow Run Connected & Autonomous Vehicle Testbed
- MDISS World Health Information Security Testing Lab (WHISTL)
- Various CPS, IoT, SDN, and Cloud Cybersecurity Testbeds

What's Next? CEF Shared, Connected Ecosystem

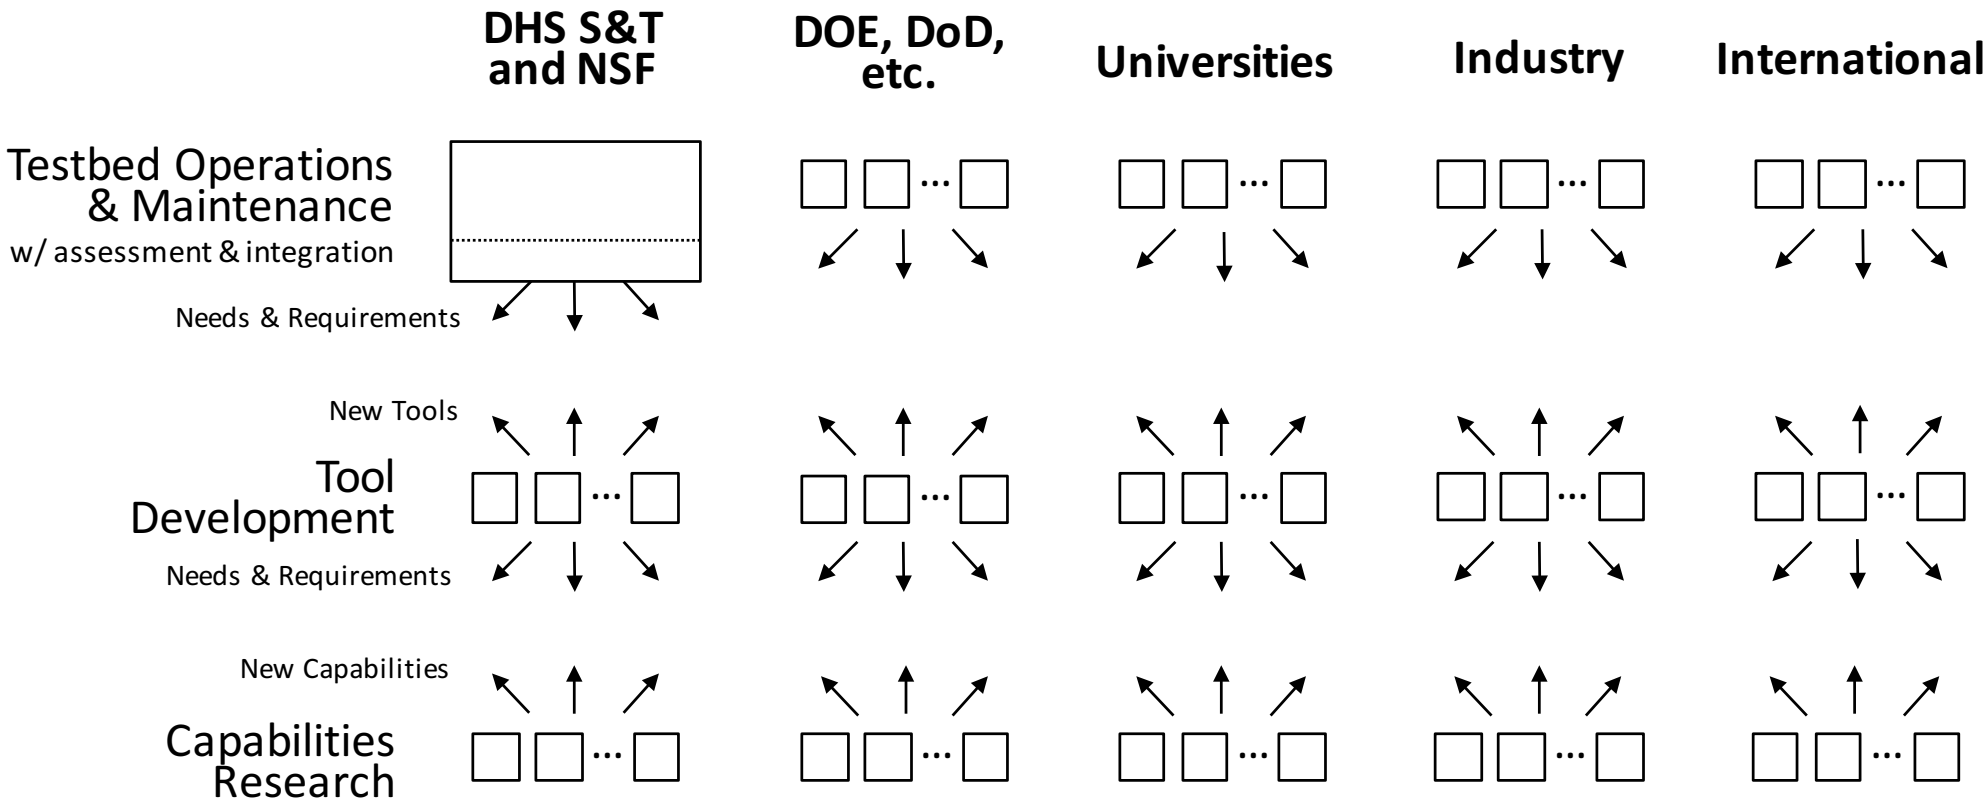

Collective needs and requirements drive collaborative development of new experimental capabilities and tools across government, academia, and industry in the U.S. and abroad

What's Next? Building CEF Community

- Catalyze collaboration and sharing
  - Set up infrastructure to enable sharing and discussion
  - Community-wide identification of existing components to share – from prior research efforts

- Moving forward
  - Encourage researchers to support, adopt, and/or contribute to CEF vision as it evolves
  - Structure research efforts to leverage and/or share newly developed infrastructure
  - Identify and encourage investment in core capabilities by research funding organizations

SRI International

USC Viterbi School of Engineering
Information Sciences Institute

Thank you!

David Balenson, SRI International
david.balenson@sri.com, 703-247-8551

Laura Tinnel, SRI International
laura.tinnel@sri.com, 703-247-8533

Terry Benzel, USC Information Sciences Institute (ISI)
tbenzel@isi.edu, 310-448-9438