Final Report for Period: 10/2008 - 03/2009 **Submitted on:** 08/12/2009

Principal Investigator: Wroclawski, John . **Award ID:** 0751027

Organization: U of Southern California

Submitted By:

Wroclawski, John - Principal Investigator

Title:

SGER: National Malware Collaboratory Investigation (NMCI)

Project Participants

Senior Personnel

Name: Wroclawski, John

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Benzel, Terry

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Faber, Ted

Worked for more than 160 Hours: Yes

Contribution to Project:

Worked on technical aspects of the project (cyberinfrastructure architecture and design) supported by DHS funding for the DETER testbed

Post-doc

Graduate Student

Name: Pingali, Venkata

Worked for more than 160 Hours: Yes

Contribution to Project:

Graduate student working with project June 2008-end, supported primarily by other NSF funding. Now hired as postdoc with DHS funding to continue related work.

Name: Bhaskara, Ganesha

Worked for more than 160 Hours: Yes

Contribution to Project:

Graduate student working with project over the past year, supported primarily with non-NSF funding.

Undergraduate Student

Technician, Programmer

Other Participant

Research Experience for Undergraduates

Organizational Partners

Other Collaborators or Contacts

As part of this work we have held a number of consultations and discussions with collaborators and contacts from industry, government, and academia. A list of these collaborators and contacts is given below. In some cases these discussions were centered on the concepts and issues surrounding the NMC itself. In other cases the collaboration was focused on a more narrow aspect, such as technical cyberinfrastructure design.

Academic and Non-Profit:
Paul Barford, University of Wisconsin
Scott Borg, US Cyber Consequences Unit
Cynthia Irvine, Naval Postgraduate School
Farnam Jahanian, University of Michigan
Richard Mathews, NCSC
Rob Ricci, University of Utah
Bill Sanders, University of Illinois
Sal Stolfo, Columbia University

Government:
Tony Sager, NSA
Will Hansen, NSA
Carl E Landwehr, IARPA
Douglas Maughan, DHS
John Monastra, ONR
Carl Landwehr, IARPA
Bridget Rogers, Sandia National Labs
Mike Van Putte, DARPA

Industry:
Peter Alor, IBM
Tom Ashoff, Sourcefire
Mary Ann Davidson, Oracle
Carrie Gates, CA
Ron Hale, ISACA
Sami Saydjari, CDA
Bret Hartman, EMC
Rick Schlichting, ATT
Steve Schwab, Sparta Inc.
Stan Stahl, Information Systems Security Association
Brain Witten, McAfee
Erik Wu, Trend Micro

Activities and Findings

Research and Education Activities:

Please see attached 'Activities and Findings' PDF file

Findings:

Please see attached 'Activities and Findings' PDF file

Training and Development:

Graduate student Mr. (now Dr.) Venkata Pingali participated in this project over its last year. Dr. Pingali a) played a key role in the development and interpretation of the survey described in ths subsection; 'Findings on Existing Social Infrastructure' b) played a significant role in our overall analysis of community building activities, and c) played and continues to play a lead role in the development of the Open Discovery Network and Open

Solutions Framework described in the 'community building initiative' section of our findings on future direction and plans.

Dr. Pingali has been hired by USC/ISI as a postdoctoral researcher to pursue this research.

Graduate student Mr. Ganesha Bhaskara is working with the PIs and Dr. Pingali to further these initiatives. Mr. Bhaskara has attended a number of relevant academic and industry meetings relevent to community building in the cybersecurity area, and is presently contributing to both the conceptual development of Open Discovery Network and Open Solutions Framework and to the technical design and implementation of the supporting web tools.

Outreach Activities:

Because outreach and community formation are core components of this SGER effort, rather than ancillary activities, we have described these activities in the main Activities section of this report. Please see the subsection, 'Working with community and consortia planning', in the Activities section above.

Journal Publications

Books or Other One-time Publications

T. Faber and J. Wroclawski, "A Federated Experiment Environment for Emulab-based Testbeds", (2009). Conference Proceedings, Published
Collection: Proc. TridentCom 2009 - The 5th International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities
Bibliography: Proc. IEEE TridentCom 2009, April 6-8, 2009, Washington, DC, USA, http://www.tridentcom.com

J. Wroclawski, J. Mirkovic, T. Faber, and S. Schwab, "A Two-Constraint Approach to Risky Cybersecurity Experiment Management", (2008). Conference Proceedings, Published
Collection: Proc. Sarnoff Symposium 2008
Bibliography: Proc. Sarnoff Symposium 2008, Princeton, NJ, April 2008

Web/Internet Site

URL(s):

http://fedd.isi.deterlab.net, http://www.isi.edu/deter, and http://www.deterlab.net

Description:

http://fedd.isi.deterlab.net is the URL for software and documentation of our testbed federation implementation. As described in the Activities and Findings section of this report, federation is a key technical enabler for a national-scale, shared, malware collaboratory or other cybersecurity test and evaluation facility.

www.isi.edu/deter and www.deterlab.net are the web sites for the DETER project. Key technical innovations for the NMCI project leverage the DETER facility. Aspects of the DETER facility are being extended to better serve the specific identified needs of the NMCI.

Other Specific Products

Product Type:

Presentations

Product Description:

Briefing package on the NMCI, power point slides used in multiple presentations to academia, industry and government. See list of outreach presentations above.

Sharing Information:

Via presentations at meetings and conferences, workshops and small group meetings.

Contributions

Contributions within Discipline:

NMC-CI enables transformational research by adopting a significant departure from the traditional 'research testbed' model. The CI is being designed to enable real-time streams of traffic and malware to be funneled into the facility for dynamic evaluation. We contribute specific developments to testbed technology in

- federation - the ability to build national-scale experimental facilities from distributed, independently managed facilities. Such facilities are capable of supporting complex malware and cybersecurity experiments with a) unusually large scale and b) multiparty nature, where the parties are either collaborative or adversarial.

- risky experiment management. We contribute a new paradigm for managing risky experiments, such as those involving live malware but also with some connection to the public Internet. Our contribution, called the T1/T2 model, has been implemented in proof of concept form for specific risky experiments, and may eventually lead to automatic, formally verified support for managing risk over a broad class of experimental research activities.

Contributions to Other Disciplines:

Much of our technical work is not limited to use in the cybersecurity area. Our development of scalable federation architectures for network testbeds, in particular, is applicable to a broad range of research of the type traditionally carried out in facilities such as Emulab and Planetlab.

Contributions to Human Resource Development:

As described earlier, we have engaged two graduate students, Venkata Pingali and Ganesha Bhaskara, to work directly with a broad range of academic and industry contacts to progress the consortium and community/collaboration building aspects of the NMC concept. This is highly unusual exposure for graduate student level researchers, who's work is more often limited to narrowly technical topics.

Dr. Pingali has been now hired by USC/ISI as a postdoctoral researcher to further pursue this research. In this capacity he continues to play a lead role in the development of the Open Discovery Network and Open Solutions Framework described in the 'community building initiative' section of our findings on future direction and plans, while further developing his already unusual ability to synthesize both technical and non-technical elements to create broad research contributions.

Contributions to Resources for Research and Education:

This SGER concerns itself with cyber infrastructure and the exploration of approaches to creating a transformational collaborative consortia for research, industry and government to create game changing approaches to improving the nation's defenses against cyber attack.

We have contributed in both of these fundamental areas. First in the technology area we have established the basic feasibility of constructing infrastructure capabilities that facilitate and catalyze new forms of cybersecurity research collaboration. We have specifically made contributions to the design and development of cyber infrastructure that includes facilities for large inter organization, multi-party experimentation and test that can accommodate a range of policies and information sharing and hiding strategies across widely distributed information resources.

Secondly, we have begun the socialization process to reach out to the research, government, and industry communities who can best contribute and benefit from this new cyber infrastructure and collaborative environment. In this area we find forward thinking leaders who are interested in participating and contributing to bring the vision to fruition.

Both of these contributions provide advances for research and education in NSF core communities as well as in the broader community including government organizations and industry.

Contributions Beyond Science and Engineering:

Our project aims to mobilize a transformation of industry operating models through a change from reactive, isolated approaches to a proactive, strategic approach based on true collaborative analysis. To successfully catalyze this we evaluate economic, as well as technical, factors and modularize tools and usage patterns to cleanly separate potentially collaborative activities from naturally competitive activities.

As outputs from our SGER we have proposed two specific activities related to the building of academic/industry/government cybersecurity-related collaborations.

- The Open Discovery Network is aimed at collaborative identification and understanding of high priority problems.

- The Open Solutions Network is aimed at sharing of expertise across traditional industry boundaries and the encouragement of composite, integrated solutions.

As followon to the SGER award being reported on here, we have hired a postdoctoral researcher to further develop these ideas and begin their reduction to practice.

Conference Proceedings

Benzel, T;Braden, B;Faber, T;Mirkovic, J;Schwab, S;Sollins, K;Wroclawski, J, Current Developments in DETER Cybersecurity Testbed Technology, "MAR 03-04, 2009", CATCH 2009: CYBERSECURITY APPLICATIONS AND TECHNOLOGY CONFERENCE FOR HOMELAND SECURITY, PROCEEDINGS, : 57-70 2009

Categories for which nothing is reported:

Organizational Partners

Any Journal

NSF Award 0751027: SGER: National Malware Collaboratory Investigation (NMCI): Final Report - Activities and Findings

What were your major research and education activities?

This SGER effort is focused on a series of activities intended to foster the creation of a *National Malware Collaboratory,* or NMC. The NMC concept brings together two synergistic innovations to provide order-of-magnitude improvement to the nation's defenses against cyber-malware. The first is technical: a national-scale distributed cybersecurity test and evaluation infrastructure aimed both at advance the science of experimental cybersecurity evaluation and at breaking down community barriers and catalyzing collaboration. The second is social: an academic/industry/government consortium built around the infrastructure, and charged with both leveraging and supporting it.

In the SGER period of performance we carried out three activities central to the establishment of the NMC:

1) Research, design, and design validation critical to the technical development of a national-scale cybersecurity test and evaluation infrastructure.

2) Identification, refinement and documentation of the research agenda and methodologies *enabled* by the NMC, working with research and industry communities who stand to benefit from and participate in the Collaboratory.

3) Structural planning and preliminary outreach to government, industry, and academia in order to establish the Collaboratory consortium.

The primary objective of these activities is to build a widely credible and detailed case for the NMC, sufficient to a) draw support from key users within the research community; b) support preparation and review of research plans and proposals in furtherance of key NMC elements; c) identify and evangelize key industrial participants, and d) identify and build collaborative support for the NMC among potentially interested government agencies.

It is important to acknowledge that a single SGER award cannot possibly support the full range of activities desirable to further the NMC vision. Rather, we worked to build on a number of potential support mechanisms to advance the NMC as effectively and rapidly as possible. We are continuing discussions around the NMC vision with academia, government and industry partners and potential industrial sponsors; with the intention of working with all interested parties to foster the Collaboratory. In addition, we are leveraging as much as possible the community's existing efforts; DETER, GENI, Emulab, PlanetLab, etc, and in return sharing both specific technical research results and organizational and structural lessons with other interested parties in the larger research enterprise. One potentially key customer for our work is the proposed DARPA-sponsored National Cyber Range, and we have established multiple channels for concept and technology transfer to that effort.

Cyberinfrastructure

The cyberinfrastructure at the heart of the Collaboratory effort – the NMC-CI, or NMC CyberInfrastructure – is a national-scale, distributed, shareable evolution of today's research testbeds, such as our DETER facility (http://www.deterlab.net). Design of the NMC-CI explicitly considers both technical requirements and industry-structure objectives to obtain the transformational results we seek. NMC-CI enables new classes of malware research through its dynamic and real time nature, scale, heterogeneity, and tools. Simultaneously, the NMC-CI provides the technical catalyst to foster effective collaboration in an otherwise highly competitive environment, through federation, information sharing and mutual-benefit trading.

Activities described in this subsection focused on research, design, and design validation critical to the technical structure of the NMC-CI itself. Key topics include distributed resource ownership & federation under different models with varying security policies; the management and containment of "risky" experiments and tests. The establishment of infrastructure that simultaneously enables distributed information sharing and strong isolation/containment is the catalyst for transformational forms of collaboration. This collaboration takes the form of collaborative scenario modeling, experimentation and test with some or all of the following characteristics:

- Scale. Experiments that involve complicated composite behaviors, rare event detection or emergent effects may need to be quite large and complex to be accurate or indicative.

- Multi-party nature. Most interesting cybersecurity experiments involve more than one logical or physical party, due to the adversarial nature of the problem as well as the distributed, decentralized nature of the networked systems environment.
- Risk. Cybersecurity experiments by their fundamental nature may involve significant risk if not properly contained and controlled. At the same time, these experiments may well require some degree of interaction with the larger world to be useful. Within the context of collaborations risk is also a critical factor in defining information sharing policies. Competitive parties require capabilities to share portions of their data and experimentation approaches while still retaining key intellectual advantages.

Our initial NMC cyber infrastructure provides a proof of concept for federation, securely contained facilities, and controlled interconnection to larger world. Together, these system elements enable a) local owner/operator control of *facility resources,* with the ability to allow dynamic construction of larger or emergent experiments; and b) controlled sharing of *information,* to catalyze collaborative activities across otherwise competitive players.

Federation:

Federation is bringing together, on demand, the resources, users, and local capabilities of multiple, decentralized, separately administered communities to conduct a single, shared, exercise, evaluation, or experiment. The goal of federation is to subdivide and embed a single experiment across multiple testbeds, in a way that meets the objectives, requirements and constraints of both the researcher and the testbed operators. Reasons to federate experiments include scale and realism, access to heterogeneous testbed capabilities, integration of multiple research communities, and information hiding. Of particular interest for the NMCI investigation is for integration of multiple research communities, and information hiding.

We worked synergistically with our DHS funded DETER project and with the GENI prototype and development efforts to make significant progress around questions of facility federation and resource ownership. We developed a federation architecture and an implementation of that architecture for emulab-style testbeds. We deployed prototype versions of the developed capabilities locally on the DETER testbed, at the University of Wisconsin's WAIL facility (http://wail.cs.wisc.edu) and at the University of Utah's Emulab site (http://www.emulab.net) to demonstrate the concept of a national-scale, decentralized and federated cyber test and evaluation infrastructure.

We developed and released software and documentation describing an initial version of our federation architecture, which supports federated experiments across multiple emulab-style testbeds. This material is available at http://fedd.isi.deterlab.net.

The development of this architecture and implementation is continuing with support from other sources. Current work focuses on the design of advanced access control and policy management capabilities and on the ability to federate additional key resource types, particularly high performance network resources based on the NSF-funded DRAGON architecture and software (http://dragon.east.isi.edu).

Managing risky experiments:

A second key capability needed in the NMC-CI is the ability to safely carry out and manage risky experiments. This is because:

First, experimental cybersecurity research is often *inherently* risky. An experiment may involve releasing live malware code, operating a real botnet, or creating other highly disruptive network conditions. Realistic replication of such attacks is necessary to thoroughly test detection and defense mechanisms.

Second, both fear of this risk and actual technical inability to manage such risk limit the willingness of many researchers to carry out valuable and necessary research. This is particularly true of industrial researchers, where accidental impact on the larger world may result in substantial financial and publicity consequences.

The common technical response to this requirement is to implement strict isolation capabilities within a testbed, in an attempt to ensure that no actual damage will be caused by an experiment. Depending on the testbed, containment mechanisms may range from complete disconnection from the outside world to allowing narrowly controlled console access, and include disk scrubbing before and after each experiment.

But such containment itself is highly limiting. A fully contained experiment is hard to observe, hard to establish, and hard to control, because it must be completely isolated from its environment. Similarly, full containment is hard to create with any assurance. Sneak paths, equipment failures, and design mistakes can render containment ineffective in myriad unexpected ways.

Most importantly, full containment is not very useful. Many of the most effective and useful experiments *need to* interact with the larger environment (i.e., touch the Internet), but only in carefully controlled and well-understood ways. Thus, our objective for any NMC-CI is to move from risky experiment *containment* to risky experiment *management* as a strategy.

To further this objective have developed an approach to risky experiment management based on a very simple line of reasoning:

- If the behavior of an experiment is completely unconstrained, the behavior of the host testbed must be completely constraining, because it can assume nothing about the experiment.
- But, if the behavior of the experiment is constrained in some particular and well-chosen way or ways, the behavior of the testbed can be less constraining, because the *combination* of experiment and testbed constraints together can provide the required overall assurance of safe behavior.

We call constraints on experiment behavior "T1 constraints", while the corresponding constraints on testbed behavior are called "T2 constraints". Consequently, we refer colloquially to the concept as "T1/T2" risk management.

Leveraging intellectual inputs and support from both this award and our DETER testbed facility, we have developed a proof of concept software framework for risky experiment management based on this observation. Our framework allows risky experiments to be created and managed in ways that address three separate concerns: experimenter's experiment validity goals, testbed operator's safety goals, and experimenter's information privacy goals. Our initial implementation of this framework is relatively primitive – although it is capable of supporting certain specific classes of risky experiments directly, it primarily serves as a validation of the methodology's potential. We have identified the most critical specific research directions needed to fully realize this potential and are now in the early stages of carrying out this research.

Community and Consortia Planning

With advances in the design and early development of this architecture, under synergistic R&D efforts at ISI, we began socializing the transformational capability of NMC with research and industry communities who stand to benefit from and participate in the Collaboratory.

We undertook several activities, as part of this socialization process, to define and develop the conceptual underpinnings of the consortium further, and develop a concrete plan of execution. This included holding focused conversations with key stakeholders and potential partners, participation in related community development initiatives, conducting a survey of the existing initiatives, developing a roadmap of execution, and recruitment of appropriate individuals.

We met with key representatives from academia, government and research, (see the list below) and began discussions around identification, and refinement and of the research agenda and methodologies *enabled* by the NMC. Bulk of the discussions centered around technical advances in federation and information sharing that create the environment for new cross community research agenda. Although there are large portions of the community who are looking for the opportunity to change the dynamics of the reactive attack-defend cycle, they are first and foremost oriented towards gaining competitive advantage. Yet at the same time there is a growing shift from reactive analysis to collaborative information sharing, see for example, Kaspersky Labs, www.kaspersky.com and the Anti-Spyware Coalition, www.antispywarecoalition.org, as a few examples of this movement. The list of individuals can be found in the "collaborators and contacts" section.

We co-hosted an Industry Workshop during November 2008 as part of a series of National Cyber Defense Initiative workshops. Participants in the workshop represented key industry segments of interest to NMC. The workshop agenda included a discussion of collaborative, yet controlled, multi-party R&D that will *enables new classes of malware research*, due to its dynamic and real time nature, scale, heterogeneity, and tools. We discussed how infrastructure can provide a *technical catalyst to open collaboration* in an otherwise competitive environment. We followed up the workshop with meetings to continue these discussions. Based on the feedback, we will plan to provide appropriate early capabilities and experimental environments as we move the NMC agenda forward.

In addition to hosting the NCDI workshop, we participated in several meetings including the 2009 IT Security Entrepreneurs' Forum (ITSEF) at Stanford in March 2009, DHS Infosec Technology Transition Council (DHS-ITTC) Meeting at Stanford Research Institute (SRI) in June 2009, Online Communities Unconference also in June 2009, and multiple entrepreneurship network meetings at Los Angeles and San Diego.

We took the opportunity in each of the meetings to discuss NMC as a concept, and get feedback with respect to relevance and economics of NMC. They offered suggestions, specific problems that could be pursued, and tools that could be useful in building the NMC.

Building on these conversations and activities, we identified and began a number of further activities intended to move beyond this SGER and towards the fulfillment of our original vision - a collaborative community centered around an emerging large-scale cyberinfrastructure. These future activities are discussed in the “Findings” section of this report.

What are your major findings from these activities?

Findings on Cyberinfrastructure

On the technical side we have established the feasibility of building large inter organization infrastructure. Furthermore, it appears that by building on DETER facility we can realize many of the desired attributes. Note however, that NMCI requires infrastructure specific to the NMCI goals, collaborative, multi-party, large, risk experiments. Thus the findings here go beyond the work currently being performed by complementary projects within the DETER community and our academic collaborators at Emulab and WAIL.

Secondly we have observed that as the complexity of experiments increase there is a need for better interfaces and graphical interface structures and approaches to manage the experiment and to reduce the learning curve for new users. We found that this is particularly true for multi party collaborative experimentation and test. We determined that these systems impose different requirements in context of NMCI on users interface than on complex large single party systems and traditional research community. As a specific example, of how this mix of design goals interact, we have found a challenge in exploring the relationship between information hiding (between experiments), multi party collaboration, and design and development of infrastructure where the parties don’t share global knowledge. This is a fruitful area for our continued work.

Findings on Existing Social Infrastructure

We surveyed existing initiatives as well as thought leaders regarding the design of the Collaborative. The former helped understand the nature of existing initiatives and identify specific opportunities that the NMC could pursue. The latter helped understand the context within which they can be pursued and appropriate methods. We discuss both these in greater detail in this section.

We surveyed approximately forty existing security initiatives with the objective of understanding the landscape. Our survey documented the following:

1. There are a large number of grass roots and top-down initiatives
2. The initiatives vary widely in terms of scope, structure, goal and method
3. Many of these initiatives are multi-functional and often are cross-domain
4. Most successful initiatives are driven by clear technical needs and are goal-oriented such as developing a specific tool or dealing with a specific worm
5. Most of the communities do not have explicit sponsorships or funding models, and grassroots groups keep overheads associated low by extensively using online communication and collaboration tools.
6. Coordination between the initiatives is often adhoc
7. Information and coordination gaps are not systematically discovered or addressed

8. Goals are often not achieved due to misaligned incentives of the participants relative to the need, missing players, and lack of resources
9. Goals also not achieved often due to misaligned incentives within enterprises that receive the output of the initiatives

We discussed technical infrastructure design and the creation of agenda for the collaboration community with thought leaders. They helped identify a number of legal and other challenges in information sharing, vested interests of specific market players, incentive structures for security professionals, and economics of security investment.

We reviewed and studied the security ecosystem from an economic perspective. One key to understanding this ecosystem is that "security" is not a single problem or even a single class of problem that can be "solved". Rather, security is a collection of ever evolving problems – many of which, in today's environment, are discovered only in retrospect. Further, there is little incentive for enterprises to invest in solving these problems. Less than 1% of the employee pool in large companies such as Google and Microsoft is dedicated to security issues. The budget for IT security in the enterprise space often ranges from 3-8% of total IT budget, which is a small fraction of total operating costs. The market for security research and solutions, as a result, is driven by fear, regulation, or adhoc considerations, rather than by economic incentive, and the products are not matched well with the problems. There is a shortage of well-qualified security professionals as well.

The basic challenge of the security community and NMC in particular, is to transform the security market place by altering the long-term economics of security. The change has to both cover the cost and benefits of the hackers and the enterprises.

Their input along with the results from our survey suggested the need to be flexible with respect to the structure the consortium takes. It is likely to be different in form, dynamics and economics from existing initiatives, and the right structure must be discovered and developed over time. However, some defining properties are apparent. It goal is to fundamentally alter the economics of security. Its structure will track the nature of the problem – as a collection of decentralized problem-solving initiatives. The methods will vary significantly across domains, platforms and needs. The partners who are likely to be most interested are government, entrepreneurs, enterprises and researchers. Vendors will new markets to pursue even if the old markets do not survive.

Findings on Future Direction and Plans

A consortium or similar organization such as NMC is well suited for the challenge of transforming the security market due to NMC's industry-wide scope, neutrality with respect to solution outcomes, and ability to accept long-term returns. Although a number of organizational models are possible, each with strengths and weaknesses, ISI's university-based position as a non-profit "neutral party", together with its unique capabilities and past experience in infrastructure and community development makes it a unique institution that can help realize the NMC.

To further this goal we have developed and are acting on the following future directions and plans.

1) We have recruited two individuals – a post-doctoral researcher and a doctoral candidate – to help with the design and development of the next stage of the consortium.

2) We are pursuing two specific community building initiatives based on the observations and conclusions reported above.

The **Open Discovery Network Initiative** is concerned with gathering about the specific nature of security problems and making it available to a global, potentially collaborative research and product development process and community. The Open Discovery Network Initiative is a framework for externalizing information about valuable problems as well as constraints. Democratizing access to this information reduces the unfair advantage that some players in the market have, and enables researchers and industry to design solutions that are more likely to be adopted.

The **Open Solutions Framework** is a vehicle to discover, develop, store and share tools, artifacts, information and processes that can be the used across researchers and enterprises. Examples include secure configurations, events, and designs. The aim is to provide a framework for discovery and management of shared, collaborative solutions possibly composed from many software and research elements, rather than individual point solutions for which existing platforms such as SourceForge and Google Groups are more suitable. From an economic perspective, the Open Solutions Framework reduces investment cost and risk by creating shared infrastructure that several organizations can invest in around specific problems that they have to solve.

Value is derived by the players by use of these frameworks to solve problems than selling the frameworks themselves. The frameworks are to be implemented through a combination of technical tools – web services/applications, etc – and policy/legal agreements. The implementation of these frameworks must be seen to create a level, open playing field, and not favor any one party or industry segment.

Although these are standalone initiatives, they complement existing and new research initiatives such as DETER and National Cyber Range, and industry-driven security product design and development.

3) We will continue to develop key technical elements of the NMC-CI, specifically designed to create cyberinfrastructure that a) supports large-scale, multi-party, collaborative experiments, exercises, and tests and b) is accessible to a broad community comprising academic and industrial researchers spanning relevant sub-disciplines within security and trustworthy systems, together with educators and commercial product development teams. Among these key technical elements are:

Further development of the **federation architecture** described previously to incorporate broader resource types (classes of federants); to support advanced and assured access control and policy management for data as well as experimental resources, and to support higher level user-appropriate interfaces for different user communities

Further development of the **risky experiment management architecture** to move from an initial proof of concept implementation useful for specific experiment classes to a broadly applicable and deployable architecture and implementation based on a) a fine-grain model for T1 and T2 constraints to support the widest range of experimental scenarios and b) a formal structure to reason about constraint composition to support the strongest possible level of risk management assurance.

A new technical direction referred to as **experiment templates** implements a top-down, structured approach to capturing and reusing experiments and scenarios. The key objective of this activity is to *capture expert domain knowledge* in particular areas for reuse by others in the community. Rather than expressing an experiment as a low-level collection of resources and configuration options, an experiment template captures and represents an experiment in a semantically meaningful form, including both the basic structure of the experiment and high level *invariants* designed to ensure that the experiment remains valid through reuse and reconfiguration. The template mechanism then allows the experiment to serve as a model for a class of similar experiments, and to be modified through different mechanisms that are suited to users of differing knowledge and sophistication.

4) Finally, the above list is not intended to be exhaustive. Any successful NMC will discover, develop and execute several such initiatives over a period of time with the help of the community built around NMC. In furtherance of this larger goal, our post-doctoral researcher will be working on further developing the social structure and a concrete plan for an institute to lead the formation of proposed structures and collaborations. Drawing on the past experiences of ISI as a research institute bridging academia, industrial research labs and collaboration with industry positions us well for this undertaking. We expect that within the coming year we will have a concrete set of recommendations and an established community support network for launching initiatives in this area.