COMPUTATIONAL
CYBERSECURITY IN
COMPROMISED
ENVIRONMENTS

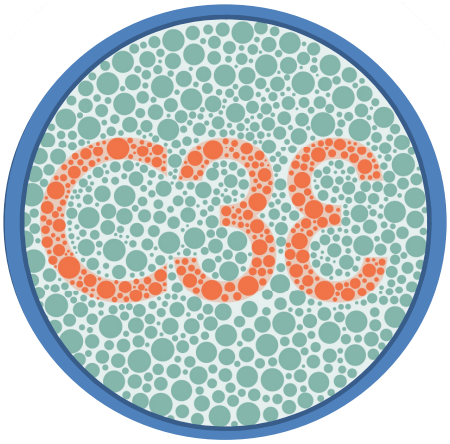

C3E Workshop Report, 2010

Contributors: Pamela Arya, Lonnie I. Carey, Jr. Kevin O'Connell, Ed Gibson, Virginia Hoover, Ted E. Senator, Alex Szalay, Daniel G. Wolf and the C3E Team

Executive Summary

The Science and Technology Lead for Cyberspace at the Office of the Director of National Intelligence (ODNI) and the Technical Director for Research at the National Security Agency (NSA) co-hosted the 2010 Computational CyberSecurity in Compromised Environments (C3E) Workshop this past August. The research workshop brought together a diverse group of top academic, commercial and government experts to examine new ways of approaching the cyber security challenges facing our Nation.

This was an analytic workshop as much as it was about cyber security. Though the problems in cyber security are many and various, and the types of expertise required to address them diverse, the group was concerned with a very specific part of the problem: how to enable smart, real-time decision making in cyberspace through both "normal" complexity and persistent adversarial behavior? The workshop was designed to identify novel ideas related to the pursuit of model, data, and human understanding of cyberspace developments.

Participants highlighted observations and findings into a set of areas for further research exploration, including:

--the need for improved models of real and potential adversaries;

--the potential for model combination and integration;

--modeling our own behavior;

--improving our understanding of cyberspace context;

--taking advantage of the comparative benefits of human and machine decision making, and

--a broader understanding of the information flows across boundaries (government, industry), especially to assist practitioners on the front lines.

While the future cyberspace environment considered here will be very challenging, we believe that it is a realistic portrayal of the future operating environment for academia, government and industry. But as our deliberations concluded, there are a number of approaches driven by models and data that can help users understand the nature of this threat, how and when it changes, and how to mitigate risk, whether at the practitioner or even public level.

These ideas are summarized in this report. As such, they are ideas that the C3E workshop participants thought to be worthy of additional U.S. government, academic, and/or private sector attention.

Introduction

The Science and Technology Lead for Cyberspace at the Office of the Director of National Intelligence (ODNI) and the Technical Director for Research at the National Security Agency (NSA) co-hosted the 2010 Computational CyberSecurity in Compromised Environments (C3E) Workshop this past August. The research workshop brought together a diverse group of top academic, commercial and government experts to examine new ways of approaching the cybersecurity challenges facing our Nation.

The workshop was an analytic workshop as much as it was about cyber security. Though the problems in cyber security are many and various, and the types of expertise required to address them diverse, the group was concerned with a very specific part of the problem: how to enable smart, real-time decision making in cyberspace through both "normal" complexity and persistent adversarial behavior? The workshop was designed to identify novel ideas related to the pursuit of model, data, and human understanding of cyberspace developments.

This was the second in a series of research workshops related to C3E. One key purpose behind these events is to create an enduring community of experts who can continue to innovate on the challenges that this unique situation of persistent adversarial behaviour creates.

Our approach this year focused on how state-of-the-art modeling activities (analytic models, systems models, adversarial behaviour models, and data-driven models) can inform the day-to-day work of the front-line practitioner, whether in government, academia, or the private sector.

The C3E Workshop explored the following areas and questions:

> --Models meet Models: to what extent do models of different behaviours inform how we think about operating in a compromised cyberspace environment? What are the emerging models that serve either directly or as metaphor for how we should think about defending and protecting ourselves in a compromised environment? Is there potential for "model mash-up" that helps inform us about how to work in this environment?
>
> --Models meet Data: how do we deal with the massive amounts of dynamically-changing data that the cyberspace environment holds, including the identification of adversarial behavior? Are the right data being collected to populate emerging models? What tools and techniques from the VLDB/XLDB communities are relevant to this problem?
>
> --Models meet Reality: how can models support real-time decision making by practitioners who deal with cyber threats every hour of every day? Are there gaps in the theoretical research that could help practitioners deal with existing, emerging, or even unanticipated problems?

Several assumptions influence our approach to addressing this problem. Cyberspace is fast moving, so decisions need be instantaneous. Cyberspace is vast and complex, so cyber analytics must work with very large data sets. Cyberspace is constantly changing, and our understanding of that space must be continuously updated. Finally, there is the realistic assumption that gives our workshop its name: cyberspace is often deeply compromised, so our analysis must be undertaken in full recognition of that adversarial dimension. These aspects formed the foundation of our work and our thinking.

Academics, government and industry officials spent three and one-half days focused on the implications of working in the compromised cyberspace environment, including the research and substantive developments that would be useful in enhancing our understanding of the operating environment and improving chances of successful and more secure computing.

During the workshop, experts first spent time in facilitated work sessions reviewing developments in the areas of models, data, and practitioner observation and experiences, in order to identify areas of possible research interest. These sessions were organized in such a way as to create synergy among people with a generally common focus in their day-to-day activities. Discussions were generally oriented around scientific and technology concepts and not around legal or policy issues.

Shifting gears, participants then looked at the issue of cyberspace analytics through the more time-pressured lens of a tabletop exercise derived from the U.S. national "Cyber Shockwave" exercise[1] that simulated a cyber attack on the United States and the subsequent U.S. response. While the original Cyber Shockwave was played at the Cabinet level, C3E participants were asked to address a series of questions surrounding analytically-based warning of the event and whether the national leadership could understand how the crisis might unfold.

The three days also included speakers tailored to a particular aspect of the C3E workshop. Participants heard briefings and lectures about the state of analysis within U.S. intelligence, modeling of human behavior in cyberspace, the view from abroad, and a series of perspectives from industry leaders about their front-line perspectives on threat identification and risk mitigation, among others.

Participants highlighted observations and findings into a set of areas for further research exploration, including:

--the need for improved models of real and potential adversaries;

--the potential for model combination and integration;

--modeling our own behavior;

--improving our understanding of cyberspace context

--taking advantage of the comparative benefits of human and machine decision making, and

--a broader understanding of the information flows across boundaries (government, industry), especially to assist practitioners on the front lines.

These ideas are summarized and organized by track theme below. As such, they are ideas that the C3E groups thought to be worthy of additional U.S. government, academic and/or private sector attention.

---

[1] Cyber ShockWave is an exercise originally created by the Chertoff Group for the Bipartisan Policy Center (BPC) in Washington, D.C. The BPC hosted the exercise in February 2010 to simulate a cyber attack on the United States and the subsequent U.S. government response. The simulation envisioned an attack that unfolds over a single day in July 2011. When the Cabinet convenes to face this crisis, 20 million of the nation's smart phones have already stopped working. The attack, the result of a malware program that had been planted in phones months earlier through a popular "March Madness" basketball bracket application, disrupts mobile service for millions. The attack escalates, shutting down an electronic energy trading platform and crippling the power grid on the Eastern seaboard.

Models

Analytical models are the engines for making rapid inferences and automating decision support needed in response to modern cyber security threats and incidents. Several leading questions were used to initiate group discussion, such as:

1) What kinds of models lend themselves best to understanding developments at the network, computer, human, or other levels?
2) What types of models best support analytical processes and tools for detecting adversarial behavior?

The main ideas generated throughout the Track and Exercise sessions that relate back to Analytical Models were:

The Value Proposition: Value vs. Risk

Currently users of cyber space and its applications have the same security environment regardless of their activities i.e. the user's banking is done on the same computer, CPU, network, browser that he/she also uses to download music or play Farmville. A new security environment is needed that models, in real time, the consequences of making security decisions and allows the user to choose between greater freedom or less risk depending on the context of their use and the "value" of their activity. The new model will be capable of feeding these separations of usage, risk and value.

Modeling the Non-Compromised System

One may best be able to recognize and detect a compromised system by comparing the state of a current system to its previous non compromised state. In order to do this, one must fully describe the uncompromised system's behaviors. The key criteria that should be modelled are: confidentiality, authentication and system integrity. In a related effort, users of any given system should also be modelled; in other words, we need to model patterns of behavior of human in cyber space. The goal of this behavior modeling is to understand many types of human behavior but at a level which eventually fingerprint individual users and entities based on their actions.

Mixed Initiative Response System

Currently there can be bottlenecks in data flow and decision making while a cyber security defense system waits for a human response. These bottlenecks can create risk and will not be a desirable characteristic of future defensive systems. We believe that a solution may be to create a mixed initiative system. This system would be created by developing a model that enable humans and computers to share in decision making based upon context and risk value. The system should enable learning and increase computer initiative by providing feedback data of the decision made and the outcome of that decision fed back into the model.

Combination, Integration and Hierarchy of Models

There are currently many different researchers, many different models, and many different practitioners with many different model preferences. It is not realistic or desirable that all these different models be collapsed into one common model or model format. Instead a methodology must be developed to integrate these models together which will not only repair their temporal inconsistencies but also answer the questions "What does this model represent? What are the boundaries that it covers?" There are additional issues related to granularity, consistency, and

confidence in and of data. We postulate that applying the theory of layer models is a good approach to solving this problem.

Continuous Targeting Strategy

To create a system to forecast what actions should be taken and to predict the consequences of these actions, we must first pre-populate the system with a set of known courses of action. The system must also input of metrics that will allow it to measure the scope of the attack.

Understanding the Normal State

Confidence in cyber warning systems can be increased over time by understanding the normal state within those systems and testing and verifying activities within it. We first need to understand and characterize the system or entity state in the absence of abnormalities. This includes modeling all kinds of users (different ages, different mental states, and different physical states). Though this is closely related to the idea of modeling the non-compromised system suggested above, it is different in that it assumes normal may very well include some level of compromise and that it includes a temporal dimension.

Modeling for Public Awareness – A Cyber "Weather Report"

Public awareness and education at a user level will be key to damage containment and virus control. One idea was to create a cyberdefense educational campaign organized in a way similar to that of worldwide weather. A website could be developed with colorful graphics and meaningful information and alerts that are accessible to the average consumer in cyber i.e. "What is the weather like on the Internet today? Turbulence brewing in …." The website could also be fed by system that acts as a cyber "tsunami warning". This warning system can focus on data collection gleaned from polling disparate sources/ people/institutions on the net: academia, industry, etc. The information would address not only what threats people are seeing but also the actions they are taking. Verification of these sources might be a key challenge. The research goals of this project would address both data collection and data presentation. Particularly in the presentation layer, one must address issues of making results understandable by the general public and preventing panic. Other consumer education issues that might be addressed here or elsewhere: user notification, usability of security measures, understandability of consequences by general users.

Data

While Models may very well be the engines driving analytical cyber security efforts, data represents the fuel for those engines. Among the questions considered in the Data Track were:

1) What techniques are useful for detecting harmful and/or adversarial behaviors in massive dynamic data sets?
2) What aggregates and summaries need to be computed and maintained for anomaly detection to be effective?

The main ideas generated throughout the Workshop that relate back to Data were:

Take Computation and Analysis to the Data

Rather than collect data centrally to learn models (of legitimate and/or suspect activity), do this locally. Potential advantages of this approach are (1) that the models can be learned more rapidly; (2) that the local models are automatically customized to unique local conditions and environments; (3)

that the local models can be applied instantaneously; (4) that there is no need for massive amounts of bandwidth and storage to transmit the local data to a central location; and (5) the potential to incorporate geo-location specific data. Global models can be learned more efficiently and equally effectively from the local models rather than from the union of all the data, and can be redistributed to be applied locally, if desirable. On-demand sharing would also be supported by this approach.

Models and Data as Two Sides of the Same Coin

Treating models as data allows models to be shared, analyzed, changed, decomposed, (re)combined, etc. In particular, treating models as data might enable a calculus of models that provides an increased understanding of what is happening at a particular time by bringing to bear the diverse perspectives of multiple models. This would include a principled method for comparing and combining not only their conclusions but also their assumptions and reasoning.

Massive Sets of Data are Models

As evidenced by Google and other organizations, massive amounts of data may be used directly as sophisticated models. These models are explainable to humans not in the usual way, i.e., by a description of abstract principles, but rather by a set of examples of similar situations in the data (e.g., a potential cyber attack that was identified by such a model would be described by showing similar situations evident in the data) combined with a description of the basic principles underlying the model (e.g., in the case of Google, their underlying page-rank algorithm). This approach has the advantage that the models are dynamic; i.e., they reflect the constantly evolving nature of both legitimate and suspect activity.

Use of Novel Data Sources

Rather than use a single data source such as user click stream patterns or IP addresses, use multiple data sources to create models with context. The additional novel data sources that are not well utilized in current environments include internet telemetry such as trace routes and router tables, simulations (e.g., results of red teams that periodically insert new types of simulated attacks into the real data stream), and potentially even data derived from crowd-sourcing or massive multiplayer games. A key advantage of this approach is that correlations of data that represent distinct views of the environment will likely be more discriminative than more thorough analysis of single data streams. Examples of novel data sources could include: (1) keystroke dynamics, (2) all internet telemetry, DNS data traffic, whois, traffic, trace roots, zone filters, routing tables, top level domains, control plane traffic; (3) data trails from analysts.

Streaming data techniques

Streaming data techniques maintain and update a dynamic model by a single-pass of all incoming data. These techniques enable both detection of relevant events or anomalies with respect to the existing model and updating of the model based on trends and patterns in the data stream.

Practice

Models and Data are useful only to the extent that they meaningfully assist a real analyst, practitioner or decision maker while performing their mission. Several questions initiated the Practitioner's discussions, including:

1) How do you evaluate the warning indicators your have given the diversity of threats?

2) Are the organizations that monitor alerts providing appropriate information (details, timeliness, etc.) on threat/attack activity?

The main ideas generated throughout the Track and Exercise sessions of particular interest to Practitioners were:

Situational Context

Context of the situation is an important factor to determine the appropriate indicators for detecting or predicting malicious activity. Models should apply multiple contexts to a scenario to generate multiple hypotheses, which would result in "tailored" situational indicators for the practitioners. Context includes internal (massive outage of internal corporate subnetworks), external (DDOS against the financial sector, explosives found in luggage at multiple airports), and technical indicators (counterfeit routers discovered in government networks, IT supply chain compromised by terrorists). The model should characterize what is normal for each of these indicators so the practitioner can identify abnormal activity. Based on the context, hypotheses should postulate the "attack scenario" and what indicators might detect the actual incident.

Exploit History to Find the Expert or Postulate the Perpetrators

A data model is needed that connects the idea of context and analytic judgment based on past signatures in either electronic or human data. What historic data exists? Model are needed to connect the ideas and analytic judgment to identify an analysts that might have insight into the attack (they worked this type of scenario before) or to a potential nefarious actor (this is the same modus operandi as before) doing the attack. Given a set of circumstances how do you find the expert for that type of scenario or how do you provide info on similar incident, i.e. what do we know from history? Do the details of the current situation/activity help to identify a potential perpetrator based on previous events/incidents? This capability would enable the practitioner to jump-start responses.

Metadata About the Data/Model

Data sources and models need to include metadata (data about data) that tag the confidence or threat level of the data or model information. What is the level of confidence in the data/model and how the analysts and their knowledge interpret the information? Also, need to have the ability to categorize the threat and how serious it is. The system should include a “confidence factor” for info provided. How reliable is the data, how confident are we in the model that provides the information, or how accurate are the conclusions of the analyst? A physical example might be to provide a confidence level on how reliable is the FBI informant. In a network, how confident is system of the attribution of an attacker? One might assign probabilities to data/models/threats as a measure of confidence/seriousness.

Track the Threat - a Proactive Approach

Need a model that picks up fingerprints/profiles of hackers that are doing practice attacks based on web traffic patterns, traffic volume, unusual incidents, or social engineering sources. What are the known hackers blogging about? What network activity is visible? What are they testing? Are they doing recon? Can we predict their next activity? How do we model these hackers or terrorists? Use this information to profile the hackers and become proactive in our defenses and responses.

Measure the Network Entropy / Provide a Human Understandable Cyber Dashboard

Need a data model of volumes of traffic, content, and domains to develop early indicators of unusual and unexpected behavior as a tool for practitioners. Can increases volume of traffic, unusual

types of traffic, or changes in Email patterns enable early detection of anomalous behavior. Need to provide a pictorial Cyber Dashboard to the practitioner to characterize the environment. What's normal? What conditions trigger an alert?

Dynamic Model of the Cyberspace

In addition to the Cyber Dashboard, need a constantly updated model of networks fed by ubiquitous sensors to understand the current state of the network and allow the practitioner to do a "what if." This model needs to include sensor data (volumes and type of traffic, loading, anomalies, etc.) from all the critical infrastructures (SCADA systems, financial networks, power grid, FAA airspace C3 systems, etc.) not just the traditional computer networks. To be effective, the model must capture the interdependencies of the systems. If the practitioner decides to close a port or suspend a service, what is the effect on DNS services worldwide, on the US power grid or on financial transactions traversing cyberspace? This model is a sandbox for the practitioner to test responses prior to executing. It's a tool to determine effectiveness of responses along with the ability evaluate unintended collateral damage in the Critical Infrastructure.

Develop an Automated Dynamic Preapproved Playbook of Responses

Implement preapproved responses into the network protection systems that take in data from sensors and execute responses at computer speed (less than 500 milliseconds). Begin with basic knowledge of known attacks and practitioners responses. As the practitioner gains confidence in the automated system, take the human out of the loop for routine responses. Design the automated system to learn from the practitioner's decisions/actions. Allow the practitioner to release responses to the automated system as confidence is achieved in the implementation of new or "learned" policies/decisions. This process will allow the practitioner to focus analysis on new/unusual/anomalous activities. The goal might be handle 80% of malicious behavior automatically at computer speed. Let the practitioner deal with the hard problems, not the routine, mundane tasks.

Conclusions and Next Steps

For the second consecutive year, under the sponsorship of the Office of the Director of National Intelligence and the National Security Agency, experts gathered to assess the human, data, and modeling aspects of a cyber environment that is compromised by persistent adversarial behavior. The ideas summarized here represent a starting point for continuing discussion and potential research attention in the pursuit of more secure computing in this complex environment.

While this environment will be very challenging, we believe it is a realistic portrayal of the future operating environment for academia, government and industry. As with the participants in the first C3E workshop, we have found useful the “bad neighborhood” metaphor as one which helps us think about how and when to mitigate risk under these circumstances. Within such a neighborhood, sometimes one evades the enemy, sometimes one calls attention to his behaviour to others, sometimes one bands together as a source of strength, and sometimes one defends or fights. But calculating what to do and when to do it is dependent on understanding much more than the adversary’s behaviour – it

is highly contextual. We believe that thinking about the challenges and any potential responses within this context is essential to understanding the path to secure computing in the future.

For sure, the C3E efforts summarized here demonstrate that we are not without options for doing so. Our brief investigation provided the starting point for looking at how using models, data and human behavior might form the basis for understanding the nature of the threat in cyberspace. Starting with our existing knowledge and models of adversary behaviors, we need to improve upon those models and compare them to many other activities that we might observe in cyberspace, whether of other human behavior—normal, accidental, benign or otherwise—or the dynamic context that surrounds them. Given the time urgency of developments in this area, we need to optimize the roles of humans and systems for what they do best, either by their very nature or by policy or system design.

Massive, complex data sets appear at first glance to be part of the challenge, but they can also be used as a landscape within which to find curious patterns of adversarial behavior that triggers defensive action and reaction. They also provide different lenses by which to observe human and system activities, and discover complex anomalies that provide warning. Finally, while models and data are helpful to our broader understanding of C3E, they are often sources of useful, maybe critical information for analysts, system administrators, and other practitioners on the front lines of defending critical infrastructure, government and private systems. Our thoughts on this even extended to ideas for alerting the public writ large about safety and security levels within cyberspace, or even trip-wires for moments when public users move from a safe or benign environment to one with more risk. Cyberspace users will have to increasingly consider the value of the information they put at risk, and how, as they move from entertainment sites to ones which hold medical, financial, or other important information.

As with many research workshops, the value comes not simply from the initial set of ideas put forward on a complex issue, but in the elaboration on that issue that comes from the continuing dialogue and critical assessment of ideas and approaches to mitigating or eliminating specific challenges. Going forward, our C3E website, http://www.c3e.info, will continue to provide the collaboration tools and ideas repository for contributors to share new and evolving approaches for advancing analytical cybersecurity.

We hope to gather again next year to continue the conversation.

Appendix A: C3E Workshop Agenda

MONDAY, 16 AUGUST			
Time	**Event**	**Speaker**	**Location**
5:00pm - 7:00pm	Opening Reception - Welcome - Informal gatherings for meetings and introductions	Co-Chairs and Sponsors	Vista I (First Floor)

TUESDAY, 17 AUGUST			
Time	**Event**	**Speaker**	**Location**
7:30am - 8:30am	Continental Breakfast		Gazebo (Second Floor)
8:30am - 8:45am	Introduction and Welcome	Kevin O'Connell Ed Gibson	El Cabrillo (Second Floor)
8:45am - 9:00am	Opening Remarks and Workshop Objectives	Pat Muoio	
9:00am - 9:15am	Detailed Review of Agenda	Ed Gibson	
9:15am - 10:30am	**The Rise of Analysis and the Need for New Analytic Models**	Kevin O'Connell	
10:30am - 11:00am	Break		
11:00am - 12:00pm	Lightning round introductions - 2 minutes each		
12:00pm - 1:30pm	Lunch with Guest Speaker: **Assessing the Human Dimensions of Cyberspace – Modeling Cyber Culture and Attack Adaptation**	Jesse Goldhammer	La Cantina (First Floor)
1:30pm - 2:30pm	Second Lightning Round introductions - 2 minutes each		
2:30pm - 5:00pm	**First Working Session:** In-Track Discussions: Leads will guide an "in-kind" session to review leading-edge ideas and developments that inform the workshop goals and objectives.	**Models**: Pamela Arya & Kevin O'Connell **Data**: Ted Senator & Alex Szalay **Practitioner**: Dan Wolf & Ed Gibson	Vista Rooms I/II/III (First Floor)
5:00pm - 5:30pm	**Introduction to the Bipartisan Policy Center's Cyber Shock Wave**	Blaise Misztal	
5:30pm	Workshop adjourns – time for informal gathering		

WEDNESDAY, 18 AUGUST			
Time	**Event**	**Speaker**	**Location**
7:30am - 8:30am	Continental Breakfast		Gazebo (Second Floor)
8:30am - 9:30am	**Leading Edge Cyber Assessment**	Dave Aucsmith	El Cabrillo (Second Floor)
9:30am - 11:30am	**Brief Out of Track Group Findings**: each of the three groups will summarize and brief out their findings in concert with the workshop goals and objectives. This will identify the gaps between theoretical approaches, analytic challenges, and practical solutions that exist within the group. Output: summary and template		El Cabrillo (Second Floor)
9:30am - 10:00am	**Models** Track	Pamela Arya & Kevin O'Connell	
10:00am - 10:30am	Break		
10:30am - 11:00am	**Data** Track	Ted Senator & Alex Szalay	
11:00am - 11:30am	**Practitioner** Track	Dan Wolf & Ed Gibson	
	Summary Review		
11:30am - 12:00pm	Teasing Out the Areas of Difference and Convergence	Kevin O'Connell Ed Gibson	
12:00pm - 1:30pm	Lunch with Guest Speaker: **Peering Into the Cyber Future from a Market Innovation Perspective**	Robert Rodriquez	La Cantina (First Floor)
1:30pm - 2:00pm	**Exercise Introduction**: Analysis Beyond Cyber Shockwave With the starting point of the Cyber Shockwave Phenomena, participants mix and split into three groups to work through questions designed to drive out new ways to improve analysis of emerging cyberspace threats.	Kevin O'Connell Ed Gibson	
2:00pm - 5:00pm	Exercise		
5:00pm	Exercise Concludes		
6:30pm - 8:00pm	Dinner with Guest Speaker: **Understanding Developments in Cyberspace: the View from Abroad**	Lord Erroll	La Cantina (First Floor)
8:00pm	Workshop concludes for the day – time for informal gathering		

THURSDAY, 19 AUGUST			
Time	**Event**	**Speaker**	**Location**
7:30am - 8:30am	Continental breakfast		Gazebo (Second Floor)
8:30am - 9:30am	**Ideas On CyberSecurity**	Steve Lukasik	El Cabrillo (Second Floor)
9:30am- 10:30am	**Brief out of Exercise Findings**: the three exercise groups brief out their responses to the exercise questions. Output: Summary and template briefing.		
9:30am - 9:50am		**Group 1**: Ted Senator & Ed Gibson	
9:50am - 10:10am		**Group 2**: Alex Szalay & Kevin O'Connell	
10:10am - 10:30am		**Group 3**: Dan Wolf & Pamela Arya	
10:30am - 10:45am	Break		
10:45am - 11:45am	**Looking to the Future: Alternative Futures**	Eric Haseltine	
11:45am - 12:15pm	**Synthesis**: this session will summarize the workshop, including major themes, things covered, things not fully explored. Methods for continuing the conversation and the path to next year's workshop	Kevin O'Connell	
12:15pm - 12:30pm		Boyd Livingston, C3E Workshop Co-Sponsor	
12:30pm - 12:45pm	Concluding Remarks	Pat Muoio and Co-Chairs	
12:45pm - 1:30pm	Box Lunch and depart for University of California Santa Barbara (UCSB)		
1:30pm - 4:30pm	Field Trip to UCSB		