ISSN 2186-7437

NII Shonan Meeting Report

No. 2014-3

Grid and Cloud Security: A Confluence

Barton P. Miller
Elisa Heymann
Yoshio Tanaka

March 24–27, 2014

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Grid and Cloud Security: A Confluence

Dates: March 24 - 27, 2014

Organizers

Prof. Barton Miller, University of Wisconsin, USA

Prof. Elisa Heymann, The Autonomous University of Barcelona, Spain

Dr. Yoshio Tanaka, National Institute of Advanced Industrial Science and Technology, Japan

Overview of the meeting

The security of Grid and Cloud computing environments is critical to today's cyber-infrastructure. The goal of the second edition of this seminar is to bring together a diverse community of researchers, practitioners, and developers, to leverage knowledge that spans the areas of Grid and Cloud security, industry and government and academia, theoretical and practical interests, and the scientific and business communities. This Shonan meeting will continue the Asia-US-Europe collaboration on security we started with the first edition of this seminar, and strengthen the efforts involving academia, industry, and government to bridge the above mentioned areas. The seminar will comprise both representative background presentations to set the context for discussions, working sessions to develop joint research agendas, and sessions that focus on joint problem-solving of a target issue selected during the week. The seminar report includes summaries of the presentations and summaries of the Q&A and discussions that accompanies each presentation.

1. Participants

	Title	First Name	Family Name	Affiliation
Organizer	Prof.	Barton	Miller	University of Wisconsin (US)
Organizer	Prof.	Elisa	Heymann	Universitat Autónoma de Barcelona (ES)
Organizer	Dr.	Yoshio	Tanaka	AIST (JP)
	Prof.	Atsuhiro	Goto	Institute of Information Security (IISEC) (JP)
	Dr.	Yoichi	Hirai	AIST (JP)
	Mr.	James	Kupsch	University of Wisconsin (US)
	Dr.	Andrew	Martin	University of Oxford (UK)
	Dr.	Shiho	Moriai	NICT (JP)
	Dr.	Leif	Nixon	National Supercomputer Centre (SE)
	Dr.	Ryuichi	Ogawa	NEC Corporation (JP)
	Ms.	Catherine	Redfield	SECOM Intelligent Systems Laboratory (JP)
	Dr.	Masaki	Shimaoka	SECOM Intelligent Systems Laboratory (JP)
	Prof.	Shinji	Shimojo	Osaka University (JP)
	Prof.	Yih-Kuen	Tsay	National Taiwan University (TW)
	Dr.	Naohiko	Uramoto	IBM Research – Tokyo (JP)
	Prof.	Chu-Sing	Yang	National Cheng Kung University (TW)

2. Schedule

90 minute sessions: Each talk will be 30-35 minutes, including time for questions during the presentations, plus a mini panel at the end of each session.

Monday morning, 9:00am – 9:30am: Opening

1. Welcome from Prof. Isao Echizen, Academic Committee Chair of NII Shonan Meeting
2. Self-Introduction by participants
3. Brief discussion of meeting organization and format; assignment of note takers.

Monday morning, 9:30 – 11:00m:

1. Elisa Heymann "Experiences with in-depth Vulnerability Assessment"
2. Shinji Shimojo "SDN and its role in Grid and Cloud"

Monday morning, 11:30am – 12:15pm:

1. Atsuhiro Goto "Cyber security education project: enPiT-security"

Monday afternoon, 2:30pm – 4:00pm:

1. Andrew Martin "Trusted Computing Technologies for Cloud Security"
2. Masaki Shimaoka "Security Concerns as Grid Computing Shifts Towards Critical Infrastructure"

Tuesday morning, 9am – 10:30am:

1. Barton P. Miller "The Software Assurance Marketplace, an Open Facility for Increasing Software Security"
2. James A. Kupsch "Automating the Use of Software Assurance Tools"

Tuesday morning, 11am – 12:30pm:

1. Shiho Moriai "On the Security of RSA Keys used in SSL Sever Certificates"
2. Ryuichi Ogawa "Software defined architecture for cloud – Is it secure enough?"

Tuesday afternoon, 1:30pm – 3:00pm:

1. Chu-Sing Yang "Toward Virtualized Security Experiment as a Service across Public Networks"
2. Yih-Kuen Tsay "Characterizing Malware Behaviors with 3-Valued Tree Automata"

Tuesday afternoon, 3:30pm – 5:00pm:

1. Catherine Redfield "Processing on Sensitive Distributed Data: Analysis and Applications of Order Preserving Encryption"
2. Yoichi Hirai "Verified Reversible Printer-Parsers"

Wednesday morning, 9am – 10:30am:

1. Leif Nixon "The Current Threat Landscape"
2. Naohiko Uramoto "Modular security and compliance services on Cloud"

Wednesday morning, 11am – 11:45am

1. YoshioTanaka "Is security ready for Big Data?"

Thursday morning, 9am – 10:30am:

1. Barton P. Miller Extra talk: "Random Testing with 'Fuzz': 20+ Years of Finding Bugs"
2. James A. Kupsch Follow-up demonstration of SWAMP

Thursday morning, 11am – noon: Closing Discussions

3. Session details; abstracts and Q&A

Monday, March 24th

Monday morning, 9:30am – 11:00am

(1) Elisa Heymann "Experiences with in-depth Vulnerability Assessment"

Abstract: I explain how the FPVA (First Principles Vulnerability Assessment) methodology was applied to assess the CREAM (Computing Resource Execution and Management) system. CREAM is part of the gLite middleware used in the largest European Grid project EGI (European Grid Infrastructure). I describe the collection of five different vulnerabilities that we found.

Summary & Q&A

[minutes taken by Bart]

Heymann presented their experiences with doing in-depth assessments and finding serious vulnerabilities in the Condor job scheduling system. Heymann's group assesses middleware; trains developers, system administrators and managers; and does research into techniques for in-depth vulnerability assessment. She presented a summary of their First Principles Vulnerability Assessment methodology, and presented their experiences assessing the CREAM unified interface for submitting Grid jobs. CREAM is used by the European Grid Initiative (EGI). In their assessment, they found five vulnerabilities. In the first one, they were able to replace user's certificates with the attacker's certificates so that the attacker could control the user's jobs. In the second vulnerability, they used a SQL injection attack based on the user using prepared statements but not using them correctly. As a result, they were able to saturate the query interface, causing a denial of service on the server. In third vulnerability, they were able to use another SQL injection to cancel any job in the system. The fourth vulnerability involved stealing the contents of the database by writing their own custom client program.

Q: Do you find that it gets harder to find vulnerabilities as you work longer on a piece of software?
A: It depends on the software. If the software was designed by a group with security experience and a strong security process, the answer is often (not always) yes. However, for most groups, this isn't the case. And new releases of software will often introduce new programs, starting the pattern again.

Q: You are focusing on software vulnerabilities. Have you looked at configuration errors too?
A: If the system is misconfigured wrong, we don't necessarily report it as a vulnerability. We will report it if is a common mistake. However we don't report every parameter that can be wrong.

Q: Is CREAM the name of your framework?
A: No, it is the name of the software system that we assessed.

Q: You use what tools are available to your group. What the main underlying characteristics that you need to capture that allows you to do these analyses in a more uniform?
A: We are able to find new kind of vulnerabilities because we look at the places in the code that are most valuable.

(2) Shinji Shimojo “SDN and its role in Grid and Cloud”
Abstract: SDN (Software Defined Networks) are a refinement of network functions and configurations. By introduction of centralized controller programmed by software, it gives you a new way to make an integrated system configured by software. In this talk, we introduce the current status of SDN technology and use cases conducted on the JGN-X, national future internet test bed. We also discussed the impact of SDN on security, security of SDN and security by SDN. Also, discuss how SDN’s are related to a new Cyber Physical Systems. With SDNs, cyber physical systems present new challenges for security and privacy.

Summary & Q&A

[minutes taken by Bart]
How can software defined networks (SDN) affect security? SDN’s defined a controlled layer over the standard Internet, virtualizing it. SDN’s allow centralized control of the access and routing for the user(s) of a particular SDN; many SDN’s can simultaneously exist. SDN’s allow you to have such control, where you couldn’t have this level of control on the underlying Internet.. Shimojo provided a description of the OpenFlow SDN switching. OpenFlow gives centralized control, network virtualization, and programmability. The security of SDN comes from centralized control. Programmability of all devices means that security is by flows not by individual routers. They created JGN-X as a test bed for next generation network development. Based on OpenFlow, they built RISE, a research infrastructure for large Scale network Experiments. With RISE, you can obtain OpenFlow-based user slices. They use SDN’s for automated malware quarantine. Malicious traffic is their system detects and reconfigures the network to isolate or block the malicious behavior. Examples of use include the Kochi Hospital successfully using SDN’s to isolate various tasks and types of information. There are new opportunities for using SDN’s in Cyber Physical Systems.

Q: How much work is it to establish an SDN? They save work but will the total work be less?
A: The SDN providers make it quite simple to install and configure.

Q: How do you know that an SDN is configured properly?

A: We may need to do new research in this area.

Q: Are there tools specifically for checking if your SDN configuration has errors?
A: You can test an SDN system with the same tools as are currently used for regular networks, though SDN's can introduce new vulnerabilities.

Q: SDN's introduce isolation. Are their mechanisms to verify that the isolation is correct?
A: Key point for SDN is the controllers. This is still an area that needs work.

Q: X509 certificates are for authorization, and for authentication you need another mechanism. We've used VOMS, and found it complicated. You have to be carefully in assigning roles. How do you verify if the role assignment and other authorizations mechanisms are appropriate?
A: You'll need to invent a mechanism for that.

Q: Do you see the security of Cyber Physical Systems as different from regular internet jobs?
A: Security requirements of Cyber Physical Systems are likely to have different issues.

Q: How do you guarantee privacy on a system where all the routes are pre-planned.
A: That's a difficult problem

Monday morning, 11:30am – 12:15pm

(1) Atsuhiro Goto "Cyber security education project: enPiT-security"

Abstract: There is serious shortage of cyber-security experts and practitioners in Japan. This shortage of cyber-security experts seems a common problem among both developed and developing countries in IT. We should make an effort to increase skilled security engineers and security practitioners ASAP, and enPiT-security is one of the education projects to solve this.

The participating universities in "enPiT-security" are IISEC, Tohoku-U, JAIST, NAIST, and Keio-U. We are providing a practical cyber-security course, named "SecCap" for graduate students. "SecCap" Curriculum covers wide-variety and up-to-date security knowledge and skills, for students in graduate schools (Master's degree) and equivalent institutes. We provide a wide range of exercises and hands-on training in areas including encryption, systems, networks, monitoring and management in the latest practical training environment. Some exercises in social engineering are provided, such as security management and business continuity management. Students choose their exercise class in support of their career target.

We provide a "SecCap certificate" to give students an incentive for their hard work, because they have to study these SecCap classes in addition to their major courses, during summer vacation and

weekends. 65 students have successfully achieved the required classes and “SecCap” certified on March 4th. We also hope that this SecCap certificate will become well known among security related industry and governments.

Summary & Q&A

[minutes taken by Catherine and Ryuichi Ogawa, slide numbers are approximate]

1. There is a shortage of well-trained security experts and practitioners in Japan. Thus NICT (National Institute of Informatics and Communications Technology) wants a program that will increase the number of skilled engineering.
2. enPIT is this education initiative. There are classes in Cloud Computing, Embedded Systems, Business Applications, and Security.
3. Outline of organizations and industry partner involved. All classes have a credit-transfer between institutions, and students at schools the main institutions have relationships with can also participate.
4. The “SecCap” Curriculum is aimed at grad students, includes practical and theoretical courses/assignments, and advanced classes.
5. Fully completing the SecCap program gives you a SecCap Certificate, since this is an additional course on top of the students normal studies for their majors.

Q (Miller): What are the prerequisites?
A (Goto): If your major is one on their approved list, you can just apply.
Q (Miller): What sort of computer background is expected?
A (Goto): Not much -- our accepted majors include social sciences and policy majors. The basic computer class gives the necessary background.
Q (Ogawa): Do you have career targets for your courses?
A (Goto): Not really. We offer our one course, and students choose projects and subjects.

6. Compulsory: Information Security Basics

Q (Tsay): What is the amount of work for each class?
A (Goto): 1.5 hours/week, for 15 weeks

7. There are advanced classes in theory, technology, and social aspects.
8. Practical exercises in the same, although many of them are technology-focused.
9. (Examples) Actual practitioners teach classes and lead exercises, and can sometimes provide real-

life data for students to practice on. Examples for forensics team projects; CTF exercises.

Q (Miller): Does your curriculum include mobile platforms?
A (Goto): Not yet.

10. SecCap 2013 (first year) had 90 students, 65 of whom were certified.
11. Expansion. Hopefully, the program will expand to more Japanese universities, companies, international collaboration.

Q (Redfield): Do you provide any of your classes/exercises to the public using an open courseware model?
A (Goto): No. The basic course allows students to access information from anywhere, but the more advanced classes and practical exercises need to be face-to-face and inaccessible because they include real-life data, and possibly sensitive information.

Q (Shimojo): Do your students get higher salaries/better jobs?
A (Goto): This is the goal. We hope our certification will carry weight. Since industry partners are involved with the training, we think that those companies, at least, will have a lot of respect for certified people.
A (Hirai): I actually took this course. For me the networking aspect was most important. I knew some people at my company when I started work, which was nice.

Q (Miller): What about continuing education? Are you planning to do refresher courses for certified people?
A (Goto): Yes. Actually our current students also included working engineers.

Q (Tanaka): Do you plan to work with technical colleges?
A (Goto): Yes, Sendai is already involved.

Q (Yang): How many courses do you have?
A (Goto): It's only one course, but students choose areas of specialization.

Q (Tanaka): How many certificates have you issued?
A (Goto): Last year was our first year, so 65.

Q (Tanaka): There's quite a distance to get to the needed 80K.

A (Goto): Yes. We're hoping to certify 100-200 next year.

Q (Ogawa): Do you include communication in English?
A (Goto): Not currently, but we think their major courses should cover that. Emergency communication and reports are covered.

Q (Shimaoka): Does this program plan to have courses for long-term industry people, too?
A (Goto): Yes, we should, especially given the number of unskilled people in the workforce.

Q (Goto): A general question: Here in Japan, engineers can only turn into managers; you can't move up in an organization as an engineer.
A (Miller): It really depends. Silicon Valley is good at having long-term career paths for engineers.

Q: I saw an ad on television for a high school security camp. Does that have any relation to your program?
A (Goto): Not exactly, but I know the camp you're talking to, and we have a lot of overlap in instructors.

Monday afternoon, 2:30pm – 4:00pm

(1) Andrew Martin "Trusted Computing Technologies for Cloud Security"

Abstract: Modern hardware and software designs are intended to allow the creation of trusted execution environments, wherein only known 'good' software is allowed to execute (or 'bad' software if present cannot hide from detection). We discussed how these technologies can be used to give guarantees of good execution to the users of a grid or cloud system. Two main options exist - one is to distrust the cloud provider entirely, and so to require evidence of all software running on individual nodes; the other is to expect the cloud provider to use trusted computing to collect and present evidence of being in a good state. The former has a minimal trusted computing base, but breaks the core abstractions of cloud computing. An additional trusted third party can help to resolve the tension between the two approaches. The evidence collected in any of these ways can be used to build strong evidence for data provenance.

.

Q&A

[minutes taken by Leif and Atsuhiro]

Q: JSM is in the user domain, so how JSM checks hypervisor? Will the cloud service provider permit

this kind of check?
A: There is a possible conflict here; the necessary fine-grained attestation may break cloud model.

Q: How do you provide the security policy?
A: This is an implementation question.

Q: Is the goal for the user to be able to trust the cloud?
A: Yes.

Q: Suppose that a user wants to use the Cloud and before using it, how can the user do the attestation?
A: The user needs to ask a third party to accomplish this.

(2) Masaki Shimaoka "Security Concerns as Grid Computing Shifts towards Critical Infrastructure"

Abstract: Many kinds of Grid technology will be critical infrastructure in the future. This presentation discusses what we should do for such a future. The motivation of this presentation is sharing the lessons that we learned from recent PKI incidents. PKI is already cyber infrastructure for trustworthiness in the Internet, but recently was the target of more serious attacks.

Some of the lessons that we learned from PKI are: (1) assuming some migration for security, because a technology such as cyber infrastructure is becoming an increasing target in more serious attacks; and (2) there are some characteristic problems: loss of migration agility, difficulty of consensus of migration, and dilemma of something in migration.

Q&A

[minutes taken by Leif and Atsuhiro]
Q: What will be the biggest change for PKI in the future?
A: There will be new PKI alternatives. In the near future, these will have similar style to PKI Further down the road, these will be very different from PKI.
Q: Considering the new alternatives for PKI, will they reduce migration costs?
A: No.
Q: Will you be able to solve the certificate life time issue?
A: Long term signature techniques have solved it for the immediate future (3-5 years). However, for the further future, this problem is not solved.
Q: Do you have estimates for the migration cost?
A: HSM replace ($3 - 60k) + auditor ( ) → 1M US$ for each migration of crypto algorithm.

Migration time period is much more critical. 7 to 8 years.

In Grid community experience, it was very touch work.

Q: What kind of security, authentication of end-user, will be appropriate for cloud? End user generally prefers “passwd” not PKI.

A: For authentication, I give up to use PKI. But PKI is useful for not auth but signature and others. ID federation and multi-factor auth will be good.

Q: FIDO, fast ID online

A: Talk later

Tuesday, March 25th

Tuesday morning, 9am – 10:30am

(1) Barton P. Miller “The Software Assurance Marketplace, an Open Facility for Increasing Software Security”

Abstract: In The Software Assurance Marketplace (SWAMP) is a new project targeted at improving the quality of open source software and the quality of open source software assessment tools. As a primary task, the SWAMP automates the ability to run software assessment tools on software packages. Users can bring new software packages to be assessed or new assessment tools to be run against software packages, and the SWAMP provides a high degree of automation for these tasks.

The SWAMP provides “continuous assurance”, providing the ability to run assessment of a software package by a suite of tools on each software commit or tool update. The SWAMP also provides a secure facility with the ability to share or keep private any given tool, software package or assessment result, as the discretion of the user. In addition, each assessment run is executes in its own virtual machine, provide strong degree of isolation.

The SWAMP is open for use, providing a suite of assessment tools for C, C++, and Java programs, and a collection of 300+ software packages. New tools and packages are being added on a regular basis. The initial tools perform static analysis, operating on source code or (for Java) byte code. Upcoming developments will include support for binary analysis tools, dynamic tools, and tools targeted at both web and mobile applications.

Users can sign up for SWAMP access at www.continuousassurance.org.

Summary & Q&A

[minutes taken by Andrew]

Software vulnerabilities -- often through uncheck bounds on strings - have been around for a long time, as has software to exploit them. There are now many other kinds of newer exploits too. Programmers need to learn to write code with security in mind - and need to have tools to help them do so. Tool warnings need to be addressed from day one – otherwise addressing their issues becomes

overwhelming later in the lifecycle.

The SWAMP is a 5-year $23M grant whose goal is to build a facility where open source software can be tested for vulnerabilities for free. The objective is to support running the tools easily - on every code update or commit. The central concept is “continuous assurance”. Key attributes: highly automated; secure (sandboxed); private (if you wish); open; resource (software and people); a community.

Open source software is widely used, and tends to be highly reliable and strongly peer reviewed: but the code is unverified, has unknown sources, supply chain dependencies, and hidden vulnerabilities. So the approach is to test and analyze open source software - aware that different tools will suit different settings and contexts.

Q: Are there really that many eyes looking at the software?
A: You should try committing some code that transgresses guidelines in some way! Yes.

Analogy with continuous integration - continuous assurance means checking after each commit.

Their audiences are: software developers, software assurance tool developers, tool researchers, educators and students, and infrastructure operators. Aiming to make this part of the everyday experience for the software developer - so providing a rich collection of capabilities for automated assessment, and results analysis. For assurance tool developers, the aim will be to improve their tool quality - reducing false positives; finding true positives - as well as to increase their user base and showcase the tool. Access to test material and statistical analysis helps the tool researcher. The platform can help in education, so those learning about software development – and/or secure development – can have access to the tool suite. There will be specific capabilities for this. For the infrastructure operator, the SWAMP is a place to understand the supply chain of software packages being used.

The SWAMP service presently has 700 cores, 5TB RAM. It incorporates a number of open source analysis tools, with more coming, and some commercial tools. It supports nine platforms - including Windows and varieties of Linux. By Q3 2014 support is expected for mobile (Android), and other tools will be added over time. www.continuousassurance.org
Q: Is there a defined level of assurance for software, using these tools?
A: People would like a ‘stamp’ rating. There is a long-term goal to provide some kind of certification of results. This is challenging: needs to produce demonstrably reproducible useful results. It also raises

complex legal issues.

Q: You support C, C++, Java. Are there plans to expand?
A: Yes; plans to expand to web and scripting languages. Also to iOS/Objective-C eventually.

Q: What are the most challenging issues for making this work?
A: The automation. There are many platforms. Bringing software in and making both it and the tools run on multiple platforms automatically is in practice difficult. Combining results is also challenging, because of the diverse reporting styles and norms. It is also interesting to consider how to provide stable results over time – as different versions of packages etc., are used.

Q: Considering the software engineering lifecycle, is there feedback to the design phase? How to integrate into project management?
A: Yes. This is part of the testing phase – as a peer of correctness testing. Ultimately plan to offer interfaces/dashboards etc., for management analysis.

Q: Can this be used as a playground for a course? E..g., for a lab where software is developed over a 16-week period, learning to use version control and issue-tracking.
A: Yes; there is a protection model to build exactly this kind of capability. Automation to make this easier for instructors is coming in a year.

Q: Is the SWAMP available for open source developers? Collaboration with OWASP, Apache?
A: Yes, we are definitely open to these groups and already talking to them.

(2) James A. Kupsch “Automating the Use of Software Assurance Tools”

Abstract: The application of software assurance tools to a software package often requires manual modifications to the build process and a different set of modifications for each tool. This talk presents the technique we developed for the SWAMP (Software Assurance Marketplace) that allows for efficient automated application of arbitrary software assurance tools to arbitrary software packages.

The talk presents how static software assurance tools operate, and what information must be known to correctly operate such a tool. It then presents our technique for collecting this information from the software package to be assessed. To collect this information we only require a description of how to build the software package. It then concludes with the current state of our work in the SWAMP.

Summary & Q&A

[minutes taken by Andrew]
The problem: apply arbitrary static software assurance tools to arbitrary software packages – with minimal manual involvement and high fidelity. There are three elements - packages, pools, and platforms: these are inputs to the assessment.

A static SwA tool is a program that analyzes source code or binary code (including bytecode). The former is the main focus, but the latter is also in scope. Results include reports of weaknesses; metrics and other facts; judgments about certain semantic issues; style/standard conformance.

In the simplest case: a simple source file is compiled (with no special options, libraries, etc.) to an executable. Another tool takes the same input, and produces analysis results instead of binary executable. However in practice, there are multiple source files, intermediate steps (creation of object files) combination into libraries and executables; build generators, custom scripts, indirection and multi-level makefiles, etc. Hence, it is not so obvious what code to assess.

The build contains the information needed, but developers often have only minimal understanding of their build system (partially declarative; partly code interpreted at build time; arbitrary programs; many layers). Overall, the challenge is one of great complexity.

The source files of interest include system include files, generated source files (after build configuration) - and the source of libraries (often missing). To operate a SwA tool, we must capture these source files, and also object files, project obj and lib files, and project tools/executables.

A typical build generates more than one executable – tools must determine which executables exist, which source contributes to the build. This is complex – and requires automation. The precise behavior depends on the design of the analysis tool – they may claim to be whole package analysis tools, but may not be in practice (or may make non-existent connections between distinct executables - “false sharing problem”).

C/C++ raises further challenges regarding command line options (e.g., search paths, language dialects). Some open source tools may make simplifying assumptions (selection of files, maybe recursive directory search, ignore command options and environment) – which may be over-simplifications, or require excessive customization.

Their approach is to monitor the build, and then apply the tool. So the user [tool provider??] must supply an operation description file to drive the SwA tool. This is a relatively simple text file. The

source provider must provide a package build description - may be a script, makefile, Ant/Maven description, etc.

The current state of the system is to support any build process for C/C++, and Ant and Maven builds for Java. The build monitor can dispatch a number of different analysis tools, and then combine the results. The assessment framework involves the build monitor that takes the build command as a parameter, runs command and records system calls, etc,. etc.

For C/C++ the framework is agnostic to the build system, and adds 5-20% overhead. It may create large intermediate files. The Java framework is specific to Ant or Maven. Binary code analysis is inherently easier to set up.

The overall costs are per package, per SwA tool, per system type, per tool application framework - a single description, etc, in each case. This appears minimal.

In the future, the system will support new tools, new versions, new platforms. Dynamic testing is a deeper/harder problem, but is in the roadmap.

Free to sign-up and try out at http://continuousassurance.org

Q: Some libraries may be provided only as binary; can these be handled?
A: Yes, in general; there are models and tools to support these, albeit with some assumptions being made about the behavior and expectations, reducing accuracy.

Q: Is the description human-readable or machine-readable?
A: It is intended to be both. Akin to the information you would supply at the command line.

Q: Does the site track user behavior?
Answer: I guess that there are logs. Maybe in the long term do research in how people do assurance. There is a privacy issue and it accessing such information has to be consistent with our privacy policy.

Tuesday morning, 11am – 12:30pm

(1) Shiho Moriai "On the Security of RSA Keys used in SSL Sever Certificates"

Abstract: Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), are cryptographic protocols which are designed to provide communication security over the Internet. The TLS/SSL are in widespread use in applications such as web browsing, electronic mail, online shopping

etc. In 2012, two teams independently reported that we are faced with a new threat to the RSA public keys that are used in X.509 certificates for SSL server authentication. The teams performed a large-scale study of RSA public keys in use on the Internet and discovered that significant numbers of keys shared common prime factors and were hence insecure due to insufficient randomness. In this talk, we present our system "XPIA", X.509 certificate Public-keys Investigation and Analysis system, which studies current status of the vulnerable hosts regarding the threat above.

Summary & Q&A

[minutes taken by Catherine]
Identifying oneself in real life is often done with a government-issued ID card, such as a driver's license. The government acts as a trusted third party who is authenticating your identity.

Translating this to the Internet, we use a Certificate Authority (a CA) that issues people and organizations digital certificates authenticating them. User A can identify user B by checking A's certificate with the CA's public key. Certificates take the standard form called X.509, and RSA is the most commonly used digital signature on a certificate.

Q (Miller): Why does RSA dominate?
A (Moriai): It's easy to implement, and in the original X.509 definition, it was mandated. Also, the patent issues that it used to have are gone, so it's safe, legally, to use.

Q (Miller): Do you think RSA is the best option?
A (Moriai): Well, we currently have other possible options, but that kind of standardization is really a global consensus.

Introduction to RSA. (wiki)
In June/July of 2012, two papers were published demonstrating that many RSA public keys were weak. Specifically, the original prime numbers chosen for the key pair (p and q) were often repeated. If two RSA public keys have the same q, you can find their GCD (which is q), and from there it is easy to compute both keys' p.

How did this happen?
→ Pseudorandom number generators (PRNGs) are using insufficient random seeds
→ Vulnerabilities in PRNGs
→ One device only generated 9 possible keys

→ Users or service provider misconfigurations

What could be a result of this weakness? When DigiNotar's CA was attacked, the CA key was stolen, and the hackers issued >500 rogue SSL certificates. Eventually, all DigiNotar certificates were revoked.

Q (Tanaka): How long did it to find and fix these attacks take?
A (Moriai): A few months.

Currently, many people are not aware of this problem, and even the people who are aware of it don't know how to check whether a host is vulnerable or not. As a result, NICT performed a full analysis of publicly available SSL server keys.

XPIA: The tools they developed to perform this analysis.
Collect certificates → Extract public keys → Analyze keys → Get IP address/location of vulnerable hosts

The state of certificates currently is that almost all use RSA. Most (~70%) RSA uses are 1024-bit; ~25% is 2048-bit (we're meant to be migrating to this, but it takes time), and the remainder 512-bit (this is just unsafe).

Visualization: dynamic map of the globe showing prime number pairs between countries.

Q (Miller): Could this weakness have been engineered?
A (Moriai): That is one possibility. The US had the most weak certificates, then China.

Q (Tanaka): That "worst" for weak certificates, you mean total number, right?
A (Moriai): Yes. So a country with more SSL hosts will have more weak public keys.

Some notes on analysis: Not all SSL servers had downloadable public keys, but of those that did, 2611 were still vulnerable. However, these weren't popular sites; they were more likely to be printers or remote hosts.

Q (Shimojo): If most of the vulnerability is the fault of the devices, can you find what vendors are most vulnerable?
A (Moriai): Yes, I had a long list when I first finished, and contacted organizations in question so that they could update their keys.

Q (Shimojo): Have vulnerabilities decreased since you started running XPIA?
A (Moriai): Yes, and we mean to continue rescanning and analyzing to keep them low.

Q (Tsay): This is the first I've heard of this, even though I teach security. When was it published?
A (Moriai): Summer 2012.

Q (Tsay): Do you charge organizations when you tell them their hosts are vulnerable?
A (Moriai): No, we're a nationally-funded group, so we contact all Japanese hosts for free.

Q (Redfield): Are the results of your research publicly available.
A (Moriai): No, that would encourage attacks.

Q (Tsay): Can people ask you to test their systems?
A (Moriai): Yes, we take requests.

Q (Miller): It would be great if you could just have a website that would automatically check someone's key for them.
A (Moriai): Those tools do exist, yes.

Q (Tsay): What is the growth rate of SSL certificates/ public keys?
A (Moriai): We just started, so we don't have that data yet. RSA is currently migrating to RSA-2048, which should lower the possibility of random number collisions.

Q (Shimaoka): Hackers can also develop tools like XPIA, yes?
A (Moriai): Yes. It only takes 2 hours to scan all the keys.

Q (Tsay): Should we partition prime number ranges?
A: I think that would be worse.

Q (Tsay): What about certifying PRNGs?
A (Moriai): That's one of our goals, yes.

Q (Tanaka): Has there been any response from commercial groups?
A (Moriai): Not yet.

Q (Tanaka): What about Japanese residential cards?
A (Moriai): Their public keys are not verified. In fact, last year, Taiwan's ID cards were shown to have very weak keys.

Q (Shimaoka): Weren't they using a certified PNRG?
A (Moriai): Yes, it was certified, but somehow there were still problems.

(2) Ryuichi Ogawa "Software defined architecture for cloud – Is it secure enough?"

Abstract: In my talk, I discuss authorization and policy enforcement issues for cloud security. First, I introduce integrated access management technology that can automatically enforce role-based authorization rules of enterprises to different cloud platforms, including SDN controller. Also I present security challenges that SDN will face. Then I introduce cloud standardization efforts (of DMTF) to provide unified operation interface for different cloud platforms. System abstraction, virtualization and new concept of software defined architecture can help automate dynamic cloud operations much easier, but it also brings security problems such as logical/physical system mapping, troubleshooting, northbound API, forensic analysis and software testing/verification. These are all big issues, and we need to pay more attention to cloud forensics (in my opinion).

Summary & Q&A

[minutes taken by Shinji]
Ogawa-san talked about a software defined architecture and security issues on software defined networks (SDN). Integrated Access Control Management is a Role-Based Access Control on the DC. They use CIM (Common Information Model) for policy distribution on the system.

The presentation introduced ideas on how to manage SDN based systems. Next, the presentation discussed issues related to SDN security, including how to prevent attacks on SDNs and how to use SDN to increase security.

The presentation next discussed standardization, and also extensions of policy enforcement to multi-cloud environment. Software defined resource control layer is an answer for this.

Q: What does the Authentication Managers do in the system?
A: Authentication managers know the network design. IAM maps role and VM IDs.

Q: Do you have any use cases for inter-cloud excamples?

A: Hybrid of public and private cloud is a typical example.

Q: Are there any company doing that? Does NEC do that?
A: It is possible. But I cannot say when.

Q: For the security of SDN, the SDN controller has to care about three domains, northbound, southbound and management. Management of inter-cloud would seem to be a very difficult issue.
To define trust domain is difficult.
A: Yes, it is difficult.

Tuesday afternoon, 1:30pm – 3:00pm: Vulnerabilities and Intrusions

(1) Chu-Sing Yang "Toward Virtualized Security Experiment as a Service across Public Networks"

Abstract: Research and industry communities have a strong demand for facilities such as test beds for the design and test of network and security technologies. This talk presents our work on implementing large scale test beds that have been created on production networks by integrating key concept such as NetFPGA, Software Defined Network, and virtualization technologies. Combined with some novel network stitching mechanisms and a GENI-compatible control framework, we are planning to continue new developments of this network test bed. In addition, we plan to connect it with other network test beds to form a federated test bed. Our system architecture provides hooks to other middleware and helps orchestrate the operation of this virtual network. The current implementation is supported by several international research projects, including a distributed Emulab facility running on a Taiwan-based cluster and a cloud federation. This talk begins by discussing the background of our work and work that is related to our system. Two real-world case studies are used to help understand the model of our proposed system. We then describe the system requirements as well as our design and implementation details. We conclude with some performance data and summary of the lessons learned from the real deployment.

Summary & Q&A

[minutes taken by Naohiko and Masaki]
This is a project initiated by Taiwan Information Security Center (TWISC). The test bed is connected with partner research groups in Taiwan and US. Our mission is Next-generation Cloud-testbed with which to configure experiments, provide remote access to experiments, and consists of multiple planes (service, control, and data planes).

Currently test beds in Taiwan are isolated. Building a large-scale test bed over production network is

needed.

Q: Are there attacks being made to the switches?
A: No. At this moment, we are not using any OpenFlow switches.

Q: How is the user authenticated?
A: By the slice manager (SM).

Q: When the VMs are located at multiple sites, does it how is performance managed as an SLA?
A: The scheduler manages it.

Q: Can the user specify the server location?
A: Yes. The Information Center allocates them. Each site is just a broker.

Q: How much of the code for this system did you write?
A: It is a modification of Geni project.

Q: Did someone perform penetration testing for this test bed?
A: No.

Q: Are you using open source software, such as OpenStack or your own modules?
A: No OSS at this moment.

Q: What types of attacks are you expecting?
A: DDoS types, such as a DNS amplifier.

Q: Do you have a plan to expand this system to cover worldwide?
A: Yes. We plan to interconnect it with other test bed systems.

(2) Yih-Kuen Tsay "Characterizing Malware Behaviors with 3-Valued Tree Automata"

Abstract: The essentially same piece of malware may be disguised as different programs through obfuscation and encryption techniques. Malware detection based on syntactic characteristics alone is insufficient, and semantics must be taken into account. Under the assumption that the behaviors of a program can be characterized as a set of finite trees labeled with actions, we propose the notion of a 3-valued deterministic finite tree automaton (3DFT) to be a foundation for the semantic analysis of

malware.

A 3DFT may accept an input tree, reject it, or classify it as an unknown tree. An algorithm is developed to learn 3DFT, from a set of positive examples and a set of negative examples.

With the learning algorithm as a core component, a prototype malware detector has been implemented for evaluating the effectiveness of the automata-based behavioral characterization.

It is hoped that malware detection with 3-valued results would be a useful building block of more accurate malware analysis. Also, the idea of a 3-valued tree automaton seems to be new and may be used for other purposes.

Summary & Q&A

[minutes taken by Naohiko and Masaki]

Background: The same malware appears as different programs by the use of obfuscation and encryption. As a result, syntax-based detection is insufficient. Naive malware detectors produce true negatives and false positives.

Our objective is to develop an automata-based characterization of program behaviors, classifying them as malicious, benign, or unknown. We are working to devise an automaton learning algorithm to obtain the semantic characterization of a malicious program, and then design and implement a prototype malware detector to experiment with the characterization.

Q: What kinds of learning models or frameworks have you used?

A: An active learning model with L*.

Our assumption is that a program performs some action in each step, and a program is equated with a set of finite trees.

Q: Is it tree or graph?

A: A graph can be represented by a set of trees.

Q: How do you map a program to trees?

A: From program trace records, we calculate data dependencies

Q: Is it done by traces or using static code analysis?

A: Other teams are working with dynamic analysis. Static code analysis is too hard for malware detection.

Q: How is the performance?

A: So far, we operated at just a small scale. It takes several hours for learning. It also depends on accuracy.

Tuesday afternoon, 3:30pm – 5:00pm: Policies and Experience

(1) Catherine Redfield "Processing on Sensitive Distributed Data: Analysis and Applications of Order Preserving Encryption"

Abstract: Sensitive data, or data stored with an untrusted third party, cannot be stored unmodified; ideally we want to completely encrypt our database without providing it with any of the encryption keys. However, to efficiently interact with the database in real-time, the database must be able to optimize and plan as usual. However, indexing tables and optimizing queries requires the database to have information about the data it is storing. Generally the required information is data order (and thus equality), so Order Preserving Encryption (OPE) is a possible solution. There are several existing OPE schemes, each with different security restrictions and guarantees. Boldyreva et al. present a pseudorandom order preserving function (POPF) scheme, mapping from a small message-space to a much larger one. However, ideal order preserving functions can leak up to half the plaintext bits. Popa et al. present a different scheme using a stateful encryption function with mutable ciphertext, which achieves a slightly modified guarantee of IND-OCPA security.

Summary & Q&A

[minutes taken by Yih-Kuen and Jim]

Homomorphic encryption, which allows operations directly on encrypted data, is appealing, but very slow (with the currently available schemes). Simpler encryption schemes with limited capabilities may still be useful for manipulating encrypted data in the cloud. Several Order-Preserving Encryption (OPE) schemes were investigated and evaluated in this presentation. Such a scheme may be useful, for example, in range queries on encrypted databases.

Some companies that use OPE include SAP, CryptDB, and Monomi.

In order-preserving encryption (OPE), cipher texts maintain plain text ordering.

What is security? (1) leaking no information, (2) indistinguishability against chosen plaintext attack (IND-CPA), and (3) pseudorandom behavior (indistinguishable from an ideal random function).

IND-OCPA (Indistinguishability of Ordered CPA): one way to achieve this is use randomly long bit sequences between the encrypted bits.

Q: What is the size of the cypher space?
A: Bounded to be no more than exponential of the size of the message space.

Q: On what granularity do these functions operate?
A: Each message is an integer or date or other type converted to integer.

Q: Does mOPE store cypher text in a B-tree format?
A: Yes.

Q: Can you see the order in the B-tree?
A: Yes, that is the point.

Q:Are these techniques available in products?
A: Yes, e.g., in cryptDB, which is open source.

Q: What is the performance overhead?
A: 14-26%

Q: What is Oracle doing in this area?
A: I don't know the details, though there was a talk by them at NII. SAP has said that their DB uses OPE.

Q: What is a use case for OPE?
A: You could run a web app with sensitive information, and safely move it to a cloud DB. If you were worried about the cloud provider or government getting access to your data, encrypt it before sending to cloud, and there no need to trust cloud provider. You can still get acceptable performance using OPE.

Q: Does more information leak if you monitor the queries instead of the data?
A: Yes, the largest amount of information leaked comes from hacking the network instead of the DB.

Q: Do you know of any uses of OPE in medical databases?

A: No, I don't know of any.

Q: When working on big data, with big blocks, can OPE be used?
A: Don't know, but it should work, but message space may get large.

(2) Yoichi Hirai “Verified Reversible Printer-Parsers”

Abstract: We describe how to develop printer-parser pairs, which are guaranteed to reverse, in a modular manner. Many of the know vulnerabilities are caused by mistakes in parsing of byte sequences. We used a proof checker Coq to make sure that when data is printed and parsed back, it yields the original data. The proof checker requires manual proving efforts, but we designed a library to allow modular development with code/proof reuse.

Summary & Q&A

[minutes taken by Yih-Kuen and Jim Kupsch]
When a piece of data is printed and parsed, the process should yield the original. The main contribution of this work is to develop, using Coq, such a reversible printer-parser pair with guaranteed correctness (code with a proof).

Other names for printer/parser include: marshalling/unmarshalling, encoding/decoding, and serialization/deserialization.

The relevance to security is that some serious security vulnerabilities have been caused by parsing errors: for example, see CVE-2011-3046 (crash) and CVE-2011-3026.

In general, it is undecidable whether a printer-parser pair is reversible. Proof by reduction to whether a context-free grammar accepts any string.

Q: What is your definition of a failure? (1) Causing something to happen inside the scope of the language or (2) producing a wrong data structure?
A: The parser can choose to return data or not. There can be a parser that reads a number, but if given not a number, it can return nothing.

Q: Can you apply your technique to a pair of parser/printers and show they match?
A: No, we have given up on verifying existing print/parsers. However we can construct a pair that provably matches.

Q: What is the performance of the constructed printer/parser?
A: We haven't benchmarked them; it depends on the grammar, but should be fast.

Q: How long would it take to create printer/parser for binary tree?
A: 3 days.

Wednesday, October 17th

Wednesday morning, 9am – 10:30am

(1) Leif Nixon, National, National Supercomputer Centre (Sweden), "The Current Threat Landscape"

Abstract: In this presentation, we examine large scale incidents in the cloud and the grid, specifically Operation Windigo. Experience shows that real world incidents tend to not target the actual grid and cloud aspects of the systems. Instead the attackers steal ordinary ssh credentials and/or compromise the underlying administrative systems of cloud providers.

It appears that attackers are starting to pay more attention to Linux servers. At the same time, the cloud paradigm means that more and more system administration is being performed by developers and end users, rather than by professional system administrators. These are worrying trends.

We conclude that more effort needs to be spent on training system maintainers in basic and medium-level security skills, that new methodologies for incident response and forensics in a cloud environment need to be developed, and, not surprisingly, that basic password authentication is becoming ever more obsolete.

Summary & Q&A

[minutes taken by Elisa Heymann]
The Nordic e-infrastructure collaboration is a distributed organization that supports development and operation of e-infrastructure solutions. Leif coordinates the security activities, and does incident response for EGI (European Grid Infrastructure).

Operation Windigo. Starting point: kernel.org compromised, and an ssh trojan. It took several months to understand the trojan, which intercepts passwords and exfiltrates them by sending fake DNS requests. This attack was called Ebury. In early 2013, hundreds of compromised web servers were discovered, with a new version installed as a shared library. The servers were hosted by cloud providers. Cpanel was compromised and all the root passwords were stolen.

Cdorked. Trojan, a small patch to popular web servers: Apache, Nginx, Lighttpd. Some users are

redirected to malicious/different web sites. The target was the user's web browser. When the user's computer is compromised, more malware is installed.

Cdorked infected machines are usually also Ebury infected.

Calfbot. Detected in July 2013. It was written in obfuscated Perl, used for sending spam. It's controlled from servers infected with Ebury.

Two large cloud providers were infected.

By intercepting network traffic Leif & team were able to see them for a short while. They determined that they stole passwords. They found some hints on the geolocation of where the attacks came from, but it's an ongoing investigation not to be disclosed yet.

90% of incidents are due to stolen or weak SSH credentials; or exploiting public vulnerabilities. Mostly they are young kids doing it for fun. Recently, there are more attacks to get access to computing resource to send spam or do crypto-coin mining. However, those attacks typically are quickly discovered.

Current trend: Cyber criminals are showing an increasing interest in Linux servers. Also cloud technology is moving responsibilities from sys admins to users.

Overall, we need to train more people in forensics, and doing forensics in the cloud is quite challenging.

More than 25.000 machines with Ebury infections have been observed. Everyday Cdorked redirects 500.000 visitors to Backhole/Neutrino. These exploits have been used to get the passwords.

Q. This is a social engineering attack (the initial point of attack).
A. Yes.

Q. Where is cPanel based?
A. In the US.

Q. Is this only specific to Cloud?
A. No.

Q. Are you working with US law enforcement?
A. Yes.

Q. What problems should academia be trying to solve to make your life easier?
A. Authentication in a right way. FIDO is OK as for the general direction.

Q. Do you really care about the actual identity of a user? I think that you probably don't care, and you just want to make sure that the user is the same person each time.
A. That's identity persistence. Sometime that's enough, sometimes it's not.

Q: What about using two-factor authentication?
A: One time passwords and two-factor may help.

Q. How quick can the Linux communities respond to an attack?
A. Very slowly. If you cannot react in 24 hours, we shut down logins. 24 hours is the reaction time.

Q. When should I share the info that my site was compromised?
A. Careful with cascade attacks, so do it quickly, within 24 hours.

Q. What's the scope of information propagation inside EGI.
A. It depends.

(2) Naohiko Uramoto "Modular security and compliance services on Cloud"

Abstract: In my talk, I introduce IBM's cloud services that are spread to the stack on cloud (e.g. IaaS, PaaS, and SaaS) across network boundaries (e.g., public and private clouds). This situation produces challenges on security and compliance; how we can cover such various environments and provide enterprise level quality of security and compliance services? To resolve it, we employ a modular approach which aims to compose base services into a complex service which can fit to various platforms, allowing configuration by users. I explain the idea with some scenarios to support industry regulations such as HIPPA.

Summary & Q&A

[minutes taken by Elisa Heymann]

Security should also focus on Big Data. There are three main points: (1) monitoring and distilling (sanitizing); (2) risk prediction and defense planning; and (3) adaptive and optimized response.

Cloud security revisited. Cloud challenges are based on several key features: highly virtualized, location independence, and elasticity. The security requirements vary, depending on the cloud model of the client.

IBM's cloud strategy:

1. Open: Contribution to standards.
2. Hybrid and full stack: Provide products and services on premise, IaaS, PaaS, SaaS.
3. High Value: SLaaS (Solution as a Service). Towards industry Clouds which requires specific security compliance, as well as governance requirements.

Q: How do you handle the complexity of displaying a very large number of entries?

A: This portal is not adequate for that.

Q. What about geographic constrains? Issues like "my data cannot go outside my country"

A. In Japan there are such requirements.

Q. How does IBM do key management?

A. IBM has its own key management product.

Wednesday morning, 11am – 11:45am

(1) Yoshio Tanaka "Is Security Ready for Big Data?"

Abstract: The recent our activities on Big/Open Data include (1) to provide Landsat-8 data from GEO Grid, (2) contribution to DATA.METI, (3) database federation project for radiation monitoring in Fukushima, and (4) planning of DATA.AIST. At this moment, most data in these activities are free and open. So we could make it available in straightforward manner without worrying about security. As the next steps, however, we need to think about the security. The main objective of this talk is to discuss what should we think about on security and how we can design and implement the security framework using current technologies. In this presentation, I will provide three use cases which need the security, (1) federation of Landast-8 data with the other remote sensing data which need appropriate access control, (2) design and implementation of DATA.AIST which provides data in AIST by LOD (Linked Open Data), and (3) use of personal dosimeter for monitoring exposed dose. Then, moving to discussions to clarify the issues we need to consider.

Summary & Q&A

[minutes taken by Bart Miller]

Current recent activities at AIST on big data include: continuous processing of Landsat data (providing usable data within 2 hours of receipt from satellite), DATA.METI (free and open data from government ministry METI), federated databases for radiation monitoring, and planning for an open data side from AIST (called DATA.AIST).
Landsat processing provides an open and free query interface to the Landsat data. The data is 50m resolution.

The Open DATA METI website provides a variety of information available from the various areas covered by the ministry.

The most frequently visited database from AIST is SDBS, spectral database for organic compounds. The goal is to provide a 5-star rating system grading data from only machine readable to data labeled and connected with other data sources in the world.

Radiation Monitoring: this activity was motivated by the recent Fukushima reactor event. There are many organizations doing monitoring, including national ministries, local government, universities and companies. Sources include data from cars, UAV, and fixed monitoring posts. So the goal is to collect the data, and provide unified access from a single portal. All the data is free and open. This data includes information about land, water, and air. Unfortunately, the existing data format is just CSV. AIST's role is to provide this data, and federate it with other sources.

One of the use cases is correlating radiation data with snowfall data. A second use case is trying to determine collective radiation doses.

Everything has been free and open thus far, but here are some examples for new use cases:
1. Cases where the open Landsat data is merged with commercial satellite data.
2. Open data, but you want to restrict it to requirements from the data provider, such as a limited number of downloads per day.
3. Provide personnel radiation monitoring data, where a user can see only the data relevant to them.

Q: If I can see the radiation for a certain place, why would you want to restrict it?
A: This is restricted data so that radiation data for an individual is only visible to them.

Q: Data is categorized in three types: data owned by government, data published by government, and data that can be checked by resident. So, maybe we can distinguish access by these types.

Q: What level of data does the government want in this data? Do they want to know from whom that data was collected?
A: They want to understand the correlation between the exposed dose and result. So, they don't need the names.

Q: The government is eager to give out dosimeters, and when they ask their staff to do that, what do you do if the staff says we're not ready yet. How do you handle this?
A: We can't do everything, so need to give guidelines to the government as to what can be done. And tell them what they can collect now.

Comment: We all know about the trend in wearable devices. These are sometimes linked to your health insurance plan. Can you use the same ideas?

Comment: Another topic you might be interested in is data mining of medical data. This work might be of interest.

Thursday, October 18th

Thursday morning, 9am – 10:30am: Infrastructures

(1) Barton P. Miller, Extra talk: "Random Testing with 'Fuzz': 20+ Years of Finding Bugs"

4. Summary

This workshop was the second meeting whose goal is was to bring together different communities of researchers and practitioners. At the meeting were cloud and grid security researchers; academic, industrial, and laboratory researchers; and Asian, European and North American researchers. As such, it was an important meeting to share with each other with the problems, approaches, and solutions in each domain. And, perhaps more importantly, share the open problems in each area. Some of the researchers had already met at the previous Shonan meeting, and several were at this meeting for the first time. It is clear that all the researchers valued the opportunity to have the time to have in depth discussions and establish the personal contacts on which to build future collaborations.

The continuity of these meetings provides a strong foundation for collaborations. Researchers are sharing ideas and techniques. Experimentalists are thinking about new theoretical topics in cryptography. Theoreticians have been getting feedback from experimentalists as to the usefulness of their ideas. Academics are exposed to industry real-world problems, and industrial attendees are hearing about the latest wild ideas. Discussions continue long into the evening, allowing the attendees to truly learn from each other. Many attendees went home with a list of papers that they wanted to read

to learn more about the work that was presented.

There was uniform agreement from the attendees that their horizons were broadened and they benefited from this meeting.