BREAKOUT GROUP B4 – SESSION NOTES

Simulation and Modeling at the Exascale for Energy, Ecological Sustainability, and Global Security
Oak Ridge National Laboratory – Town Hall Meeting
May 18, 2007

Breakout Group B4: Develop tools and methods to protect the distribution information technology infrastructure: ensuring network security, preventing disruption of our communications infrastructure, and defending distributed systems against attacks.

ATTENDEES

Ann Baker, ORNL
Walter Dykas, ORNL
Mike Fisk, LANL
Ian Foster, ANL
Victor Hazlewood, ORNL
Greg Hinkel, ORNL
Louis Howell, LLNL
Virginia Myrick, ORNL
Mildred Ofosu, Morgan State University
Doug Olson, LBL
Mike Papka, ANL
James Peery, Sandia National Laboratories
Don Petravick, Fermilab
Thomas Potok, ORNL/UT
Stephen Racunas, Stanford University
Nageswara (Nagi) Rao, ORNL
Frederick Sheldon, ORNL
Matt Smith, ORNL
Scott Studham, ORNL
David Vasil, NCCS/LCF
Becky Verastagui, ORNL

PRESENTATIONS

Protecting the Information Technology Infrastructure, E3SGS Town Hall Meeting, Cybersecurity Breakout
Scott Studham (Chief Information Officer, Oak Ridge National Laboratory)

(Cybersecurity Analysis – Intelligent Agents)
Thomas E. Potok (Applied Software Engineering Research Group, Oak Ridge National Laboratory)

Security at Scale
David Vasil (High Performance Computing Operations, Oak Ridge National Laboratory)

"Town Hall" Meeting, DOE Exascale Computing
Greg Hinkel (Computer and Network Security, Oak Ridge National Laboratory)

Misuse-Tolerant Distributed Systems
Mike Fisk (Computing, Telecommunications, and Networking Division, Los Alamos National Laboratory)

B4 in E3SGS: Open Science, Cyberinfrastructure Facilities, and Distributed Systems

Doug Olson (Computational Research Division, Lawrence Berkeley Laboratory)

Protecting the Information Technology Infrastructure, E3SGS Town Hall Meeting, Cybersecurity Breakout
Scott Studham (Chief Information Officer, Oak Ridge National Laboratory)

Scott Studham: (Introductory presentation to set agenda and format, followed by introduction of the panel members by Victor Hazlewood.)

DISCUSSION

Nageswara (Nagi) Rao: A third point to be considered is to build networks with dedicated channels to open control planes/user planes to users; Networks of dedicated connections, no firewalls or IP filters on dedicated connections; Security aspect less stringent.

(Cybersecurity Analysis – Intelligent Agents)
Thomas E. Potok (Applied Software Engineering Research Group, Oak Ridge National Laboratory)

Thomas Potok: Discussion of work with intelligent agents running from system to system. Have split seconds to make decisions with limited information. Today there is a stack of paper vs. stack of DVDs. Main focus of R&D is to get through the information. Snort packets and flagging packets – large volume of information (1M hits/day) – how to use it? Key technology is agent technology, P2P topology, blackboard coordination/control and message encapsulation. Use cluster computer infrastructure and Java (better in cluster).

Look at data in terms of packets (for cyber analogy). The process today is that we get data quickly (minutes), analyze the data (months), then review the analysis (minutes). Most of the time is in the analysis. Piranha can be used to organize raw data to categories that are linked and organized together to reduce time for analysis from months to days. We can set frequency over a large dataset for analysis as opposed to looking at small sets as they come in. It is based on words and frequencies – apply to tag documents and words. Tested versus term frequency (term weighting schemes), test data sets, and cluster algorithms resulting in high results (see bar graph) showing that it works.

Can have documents/packets information coming in, convert to vectors, and compare rapidly to see organization of data (days to now minutes).

Current cyber security approach, for huge attacks it is easy to identify the anomaly or able to repel known attack signature. But what if the attack is unknown and unusual? To find it you can identify potential targets, find machine/scientist and identify who may target that machine – can use this to characterize people behavior. Take Snort (filter

for packets – 99.9% false positive), then take IP address – go in DB NetReg, link to proposal database, link to suspect database. Get clusters to identify these suspect Snort flags, then based on this (looking at scientists profiles), identify attempts to log in to a targets system (a forensic approach).

DISCUSSION:

Scott Studham: Have you analyzed logs for cross-sites?

Thomas Pottock: No just packet headers for ORNL data – hard to get data, so only look at headers

Don Petravick: What about sys log, app log – would you expect the zip logs to compare?

Thomas Pottock: No structured text words, can chara taxonomy to use to id frequency, can share some tools

Don Petravick: Is this a general problem at the exascale?

Thomas Pottock: More military guys, nuclear concern regarding targeting science; don't know if problem on exascale for R&D, but if can identify targeted scientists, may be able to identify attacks.

Don Petravick: What is the threat is the question?

Victor Hazlewood: Intelligent agents look at small scope – future is looking at aggregate netflow, sys log, enterprise, then into distributed computing environment. It is hard to get sharing between sites of log files.

Don Petravick: What about service based computing – people surprised by the authentications per day at Fermilab (~400K authentications/day because batch oriented and move a lot of data)

Victor Hazlewood: What is suspect behavior for user?

Don Petravick: Can't really identify threats because we don't know what they are – crosssite analytical environment – local single sign-on – fear it is a huge scalable data problem to do analytics across labs.

Thomas Pottock: More energy in if we can exploit what we have instead of looking for more data – would rather have less stuff and the right stuff vs. someone overtly stealing what we have.

Frederick Sheldon: Less performance to higher performance – what kind of lift did give analysis, can you extrapolate to exascale?

Thomas Pottock: Cluster helps, but the big challenge is messaging – with agent system anyone can talk to anyone. The curse is that messaging can grow (not linear growth). The benefit is mostly from algorithm as opposed to hardware; better messaging – swarming, artificial life.

Security at Scale

David Vasil (High Performance Computing Operations, Oak Ridge National Laboratory)

David Vasil: Mixed workloads running on same systems, want to remove information remnants from memory; system partitioning may be important at the exascale. Perform checksums on petabytes of data; invalidate large portion of research if not. Need systems that allow quick verification of large complex distributed systems to ensure data is not compromised and systems are operating at optimum performance. Need redundancy for large complex system; need to operate with several faults at a given time. Need to contain incidents on hosts and accounts and quickly react (event correlation, aggregate of logs and data sources). Need to maintain data utility on the system.

Science vs. security - Trade-off in productivity is a problem with science (data efficiency vs. security – may have 10% just for security). With science it is important to share information, but also need to restrict flow of information, from those with no need to know. Need interconnection at high speeds, but restrict connections to ensure confidentiality. Bleeding edge required to provide performance, but problems for regression testing needs to be done for security flaws. Unattended authentication (functionality good batch job), but need strong authentications. Users from outside orgs using systems, users don't consider security – opening system to many thousand users, the probability for one who is careless is high, need to tightly control it; knowing who is on system at any given time.

With regard to future concerns, need low maintenance, but difficult to exploit and stable (low and hardened). Run heuristics to determine anomaly of users. For distributed checksums it would be nice to use cluster file system and cluster at same time to analyze segments of data and put in single checksum. With regard to event correlation, need to use more data in aggregate and act quickly.

DISCUSSION:

Don Petravick: The problem is how default responsibilities are assigned – even a big DOE site cannot totally influence software big collaboration takes – lack policies (policy thing hard). Life would be good if someone came in and were in a sandbox which has a facility which has a perimeter – we maintain the perimeter, but the group has responsibility for not leaking the info (distributed responsibility). It is hopeless for us to control quality of all this – anything could be in the code.

David Vasil: A flaw because lots of systems have not planned for security. Industrial model – hosting facility – squarely only infrastructure for tenants, but web service up to you and not their responsibility

Don Petravick: We cannot say it is not our responsibility to an experimental group, but great progress would be a standard interview for all sites to trust – question is how we are allowed to organize ourselves – if we secure systems.

Victor Hazlewood: Good point – more complicated when DOE wants sensitive runs, bar raised on fences between users/applications. Want to run ITAR on LCF doesn't mix with foreing researchers

Don Petravick: Useful to look at sandbox – appropriate for DOE because science on larger scale.

Victor Hazlewood: X4 control – people ask about fences

Don Petravick: First we need common vocabulary of sandbox controls. Second we need policy or operational control to standardize discussion with experimental groups, standard schema of threat models. Many science groups at sites – need a document that is standard enough, semi-standard agreement – discuss once, understand, and then put to market. If target is sandbox technology, then discuss standard items and diligence of distributed group policy discussion – these two things are important if thinking of operational security. There is a variety of threat levels, don't want stupid payload in software; with industrial collaboration may have industrial espionage need to address. Shouldn't take it to the soldiers and terrorists; nice to standardize threat schemes per level with scalable security.

Victor Hazlewood: Need information limits – what is left laying around

Don Petravick: Sandbox will allow to clean up

Scott Studham: Regarding policy perspective summary – politically to sell petascale people committed to engage industry. Have low threat level to moderate threat level (Boeing, DreamWorks), so had to move to moderate controls. DOE certifies systems, talking about certifying sandboxes inside so have moderate sandbox, low sandbox.

Don Petravick: Sandbox comes with classes, overall container and sandbox engineered to highest level for any person in it, but within classes can have category of sandbox to not treat all as same.

Scott Studham: Policy for LCF is always medium now, so any low is still going at moderate because policies are at machine level and not application level.

Nageswara (Nagi) Rao: If we have dedicated networks, one machine to another, not series of firewalls

Scott Studham: Mitigate risk for point to point connections, but today the thought is to move the way we write policy to sandbox level not machine level.

Mike Fisk: The difference is outward facing piece of sandbox – because system is protected for access to sandbox, sandbox can be protected at lower level as appropriate given that we have export controls, but can we participate in open science (hard to do with very tight controls).

Stephen Racunas: Consider it an information ecosystem, people consume information and are not allow to consume some information. Can assign roles/niches to help to decide who gets what when – drafting a model of users into security model will put infrastructure security in the middle to be facilitator (not necessarily gatekeeper).

Mike Fisk: If writing a 10 year research program, the output would be white paper to drive research agenda. What needs to happen over next 10 years to get to exascale? We talk about what we are currently doing, but would encourage us to think what are the barriers to exascale in research program? The discussion is to move perimeter from outside machine to inside.

Don Petravick: Can we take this discussion and write crisply so that people at high-levels understand? It is obvious R&D program follows.

Frederick Sheldon: Want to protect and support exascale but also should think about prevention and harnessing capability.

Scott Studham: Last few years computers move to moderate (move entire computer) – but instead make policies for sandboxes. Need to figure out how to get permissions and how to do it. Nice if implementation of policy is standard across sites.

Don Petravick: Would allow distributed computing – piece for those that want one thing or another; if have to interoperate with security – want to limit number of discussions

Frederick Sheldon: Today we don't have any cyber security drivers that say we need these?? Can we use cycles to predict cyber security relevant (like information discovery – low and slow – salient – what is the data telling us)?

Louis Howell: This discussion dovetails to applications people and fault tolerance of hardware systems; similar to concern for security validations.

David Vasil: Whose problem is it if a hacker changes a few parts of the dataset? Nice if researcher had knowledge that it was changed so they do not do post processing analysis and get wrong results.

Greg Hinkel: There are three pieces of data integrity: file system, move in/out archive, and in transit (e.g., petabyte dataset – almost guaranteed that if you move it guaranteed to have issues like bitflipping). Need tripwire to check in archives to see if it

was read wrong or if a hacker (but may be tapedrive failure or associated with moving data around).

Don Petravick: In the disk drive industry, blocks have extra room for bits, so improved bit-error rates for capacity.

Mike Fisk: At these data rates, probability of bit error that still passes these checks is also non-trivial – random bit errors can plague.

David Vasil: Problem of doing checksums.

Nageswara (Nagi) Rao: Concerns are true for any scale, but DOE want us to pay attention to exascale. We have to be inclusive to state what we can do at secure levels – need something to get us to the exascale.

Ian Foster: Security problem and hardware problem regarding data integrity; also health and security is same (both security and hardware problem).

"Town Hall" Meeting, DOE Exascale Computing
Greg Hinkel (Computer and Network Security, Oak Ridge National Laboratory)

Greg Hinkel: We know that cannot prevent all intrusion. Communication is a big deal to get information to those who need it. Users need to know how to use tools properly; know how to properly share files; make users understand home computer security issue when logging in from home.

TeraGrid has a good model for incident response; build infrastructure for common community. TeraGrid communicates on small level stuff and large and has a moderator to keep conversations reasonable and sensible. DOE list doesn't have that (just maintain CIAC list). Good that lists are tightly controlled (TeraGrid is pretty tightly controlled). But need to keep management at bay up to DOE requests for compliance. TeraGrid has had some problems with cleanup – what are you going to handle when machine gets compromised depending upon type/scale of compromise.

The future needs more anomaly and less single signature based detection (most signature based, but less single); anomaly uses data from Snort with sys logs firewall logs web logs to identify the real problem. Separation of supercomputer from desktop and infrastructure – similar to sandbox – determine what type of traffic will be monitored to get more manageable volume of data (i.e., low-volume may need to be monitored more due to level of control). Consider the types of data and data separation needed to help analyst determine what is going on.

DISCUSSION:

Mike Fisk: Was TeraGrid security planning self-assembling plans?

Walter Dykas: The TeraGrid working group originally was a policy group, but when incidents occurred, it grew to a technical group.

Greg Hinkel: DOE has an incident response group (that CIAC keeps track off) – very low volume; no peer group for HPC colleagues.

Doug Olson: NERST has an informal group (kept informal because of issues associated with formal groups – bureaucratic stuff).

Mike Fisk: Informal sharing information okay, but formal information sharing is a barrier to open collaboration to peers at other sites; perhaps part of the answer is projects should specify data to collect and share with project as a whole – if have job running on systems and job collects data, that data shared with security infrastructure for the project (suggestion)

Scott Studham: With regard to separation of data transfers, Nagi has thoughts on networking with system to system data channels which may need less analysis.

Greg Hinkel: There is no difference in the analysis from terascale to exascale.

Scott Studham: More bulk data transfer for exascale and the question is how to analyze it.

Greg Hinkel: Not much analysis of that kind of traffic because it is all together/not split – but also have to tell users how to configure their clients to use certain ports.

Nageswara (Nagi) Rao: There are few protocols for bulk data transfers, some do use a specific port, but are not widely deployed. Anomalies are opposite of normal behavior specific to applications, groups or apps, etc.; may facilitate analyst to characterize behaviors to define anomalies (need to characterize normal vs. anomalies for exascale level).

David Vasil: There is a lot work on end-to-end applications, much will be low traffic – so may be a different control port.

Mike Fisk: May need to retool to deal with service-oriented traffic.

Victor Hazlewood: Some traffic may move to portals rather than computers.

David Vasil: Need portals to have better logging to deal with this.

Misuse-Tolerant Distributed Systems

Mike Fisk (Computing, Telecommunications, and Networking Division, Los Alamos National Laboratory)

Mike Fisk: So far the focus on external access control but attackers can get around these authentication themes (hijacking). How to achieve misuse tolerance? This is related to sandboxing technologies we have been talking about – restrict access to resource by process group rather than by a UID to preclude stolen identities. Need to apply well-known technologies better. To delegate credentials with minimal capabilities is important – build software to generate credentials important as well as system ability to enforce credentials

The notion of middleware environment to help collect audit logs, will make it easier to share data across sites (rather than just on internal resources). We have experience with change detection with regard to anomaly detection – see what changes from status quo to detect different behavior of users. The time of response is important – real-time automated response, single clearinghouse system to generate real-time automated responses to make systems more resilient.

Need to tolerate misuse not just assume you can avoid it – consider top-to-bottom engineering to make misuse easier to tolerate and make systems resilient.

DISCUSSION:

Thomas Potok: Is it a credible threat for someone to misuse system for nefarious reason?

Mike Fisk: Not aware of cases to date, but may be a possibility. Mainly a public perception issues for having such a huge resource compromised – important to defend against; prankster can have large impact because of the perception.

Ian Foster: Is a big machine target for service attack?

Mike Fisk: It used to be when bandwidth was scarce. It may be today, need to do anomaly detection on bandwidth usage.

Don Petravick: If we can force users/researchers to normalize, they can declare their endpoints. Conceptually it is easy to know what the likelihood of who else is in VO at the moment to detect anomalies, detect finite set of perimeters for experiments. Post experimental activities – fuzzy because no data models with regard to experiments and their behavior, so need sufficient enterprise data model against open science to know for each process what the characteristics ought to be (important job) in order to remove guessing.

Mike Fisk: In terms of sandbox, may put filters in place.

Thomas Potok: Can help state that this process doesn't match or help identify errors in codes for users.

Mike Fisk: These deployed locally within our site to reduce risk profile (e.g., managing desktops). If we assert clients are in shape, then can reduce risk. Vulnerability scans on lots of machines versus managing dumb terminals (retrolutionary). Removing control from the office for malicious insider and outsider.

B4 in E3SGS: Open Science, Cyberinfrastructure Facilities, and Distributed Systems
Doug Olson (Computational Research Division, Lawrence Berkeley Laboratory)

Doug Olson: Applications for exascale, the theme is groups doing these things are sizable collaborations of people (few 1 or 2 investigator thing), unprecedented for computational scientists. Implies need for sociological changes and additional cyber infrastructure functionalities – international scope is different as well. We need to stress that systems need to be useable to support science, to preclude workarounds by scientists. We will need to parallelize even simple tasks to go to exascale with regard to system monitoring; advocate more complete BRO-like monitoring with active feedback for network monitoring and control on exascale.

With authorization, the main model is UNIX, but no correlation between DOE facilities to set privileges in one place to be applied across sites. With collaboration, it is hard to make the case that R&D is needed, but supporting more robust infrastructure is a consideration; usability issues testbed as well as workflow descriptions to test security infrastructure prior to deployment across laboratory systems. ISP like the water company (today is dirty water) – may help performance issues crossing site boundaries in exascale world.

DISCUSSION:

Mike Fisk: We can separate it to two things, knowing what/how to filter and then applying filters. Filtering pipes for exascale will be hard, can be massively parallel and may mean separating into two nodes. Having sites filter individually is not practical. Need to collect data and come up with cohesive would be better use of resources – applying filter at ES level may be challenging.

Don Petravick: With grid – what does it mean to be authenticated? Lack a critical internal structure – don't have sandbox for small communities to contain them and not be concerned on what they do in the sandbox.

Nageswara (Nagi) Rao: Separation is blurry regarding big pipes.

Doug Olson: Big data may not need filtered at exascale.

Don Petravick: Encrypting gratuitously can work against you – know less about node.

Frederick Sheldon: Investing heavily in network monitoring may have a big payoff.

Nageswara (Nagi) Rao: It needs to be done.

Don Petravick: There is so little context if all you can do is tap the network and everything aliased through ports. Suppose VPN for all PI investigative groups, then know

purpose for flow, no encryption within VPN would give you a lot of information bound to packets than if try to obscure it – can only do a better job if you have context.

Greg Hinkel: Need to mix the monitoring

Don Petravick: If big systems are in same VO, no privacy requirement, no need to worry.

Thomas Potok: Scientists would favor such a system where no firewall stuff is needed between members of the small group.

Don Petravick: It is a small-town philosophy to build infrastructure to keep them within communities.

Thomas Potok: Makes it a better sell to say optimize here, but not here.

Stephen Racunas: We can go beyond that, you can know who is doing what, same as ecosystem view (sandbox).

Don Petravick: Sandboxes give you more context for anomalies, so you do not have to worry about what happens within sandboxes, big VOs can add security within sandbox (VO is virtual organization within the grid world – experiment).

Nageswara (Nagi) Rao: It is clear that this task will not solve itself and this is the message to communicate – we need structure in this problem space to know that we need to address it.

Don Petravick: The point is technical leadership in the security dimension. The sandbox makes security problem scalable.