Networking and Information Technology
Research and Development

***The government seeks individual input; attendees/participants may provide individual advice only*.**

Middleware and Grid Interagency Coordination (MAGIC) Meeting Minutes[1]
October 3, 2018, 12-2 pm
NCO, 490 L'Enfant Plaza, Ste. 8001
Washington, D.C. 20024

Participants (*In-Person Participants)

Joe Breen	Utah
Richard Carlson	DOE/SC
Rob Gardner	UChicago
Padma Krishnaswamy	FCC
Joyce Lee*	NCO
Bryan Lyles	ORNL
Joe Mambretti	StarLight/Northwestern
Shawn McKee	UMichigan
Sean Peisert	LBL
Rajiv Ramnath	OSU
Von Welch	IU

Proceedings
This meeting was chaired by Richard Carlson (DOE/SC). September 2018 minutes are approved for posting on the MAGIC website.

Speaker Series:
- *SLATE: A new approach for DevOps in distributed scientific computing facilities*, Rob Gardner, University of Chicago; Co PIs- Joe Breen, University of Utah and Shawn McKee, University of Michigan
- *DevOps and Software Defined networking (SDNs) and Software Defined Exchanges (SDXs)*, Joe Mambretti, Northwestern University, Metropolitan Research and Education Network, StarLight
- *Software Assurance Marketplace (SWAMP) – A Continuous Assurance Platform*, Von Welch, Indiana University

Speaker Presentations
SLATE: A new approach for DevOps in distributed scientific computing facilities - Rob Gardner

What is SLATE?: Services Layer at the Edge – about mobility and capability (Slide 3)
- SLATE provides additional capabilities to build distributed infrastructure across the larger computing landscape.

[1] Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program.

- Goal: to equip the Science DMZ (Sci DMZ) with a service orchestration platform and enable the creation of federated systems. Once in place, it becomes a DevOps infrastructure for those developing services not living within a single data facility; spans multiple provider endpoints.
- Motivation: to enable multi-institution collaborative science. (Slide 5). Examples:
    - XENON (Slide 6). International collaboration with 25 institutions addressing fundamental science questions. Helped deploy distributed data infrastructure across Europe, U.S. and Israel.
        - Infrastructure diagram: Collection of distributed data inputs with endpoints into the Sci DMZ in institutions that pledge storage capacity to the collaboration - storage, replicating and processing into different distribution infrastructures (software connecting centers in Europe and U.S.). This exemplifies what we wish to rapidly build.
    - Open Science Grid (OSG) (Slide 7): Nation's distributed High-Throughput Computing (HTC) cyberinfrastructure. Initially created and driven by need to do distributed HTC for the Large Hadron Collider (LHC), OSG now supports many science domains. Configuration is a set of services deployed on endpoints on Sci DMZs. Relies on service administrators to keep services updated, deployed and operational; cuts out fast moving evolution.
    - R&D project to explore "data lake" (CERN) (Slide 8). Due to voluminous data, LHC will need new models to deliver data and flexible infrastructure to create and deploy new systems, test, etc.
    - Caching network for IceCube & LIGO: Summer 2018 project with the Pacific Research Platform (PRP) to deploy caching servers inside the network/Internet2. Peering points are located in Kansas City, Chicago and NY. Managed in modern DevOps fashion: Kubernetes (k8s) is accessible to a DevOps engineer in San Diego who quickly deploys caching servers in different places in the network. Model going forward.

Deployment is difficult (Slide 10)
- Trying to evolve our approach towards these systems to be more forward looking: add flexibility, ability to develop, innovate and bring more capability to the DevOps engineer.
- Broken DevOps cycle for a distributed infrastructure; no coherent, self-contained system. Hence, SLATE.

SLATE Vision (Slide 12)
Diagram: High-Performance Computing (HPC) Center with an edge network with a service orchestration layer allowing for collaboration. Can place services inside the resource provider's Sci DMZ so long as they conform to security policies and privileges. Extends early principles of distributed trust. Examples:
- Xenon: automate DevOps (Slide 13)
- OSG: evolve so have a few DevOps engineers in essential team deploying them. (Slide 14)
- IceCube: Following DevOps model- uses modern tools to deploy caching servers into the network, making them available to the collaborations. (Slide 15)

SLATE is **S**ervices **L**ayer **A**t **T**he **E**dge (Slide 16)
Underlying the programmable interface to a set of cyberinfrastructure resources available for science collaboration. DevOps-friendly concept of continuous integration, rapid ease of deployments.

SLATE Concepts & Components : (Slide 17)
Base layer: container orchestration (k8s)

Top layer: Federating layer provides interface to respect policies, privileges and trust relationships. Does not arise from k8s (service orchestrations aren't multi-institution, multi- domain, collaborative platforms).
Also: developed API server, persistent data store, web portal and Command-line Interface (CLI).

Website (Slide 18) for User or DevOps engineer to access via InCommon credentials.

- If coming in as resource provider, can register the SLATE edge cluster to trigger SLATE-readiness process. Can also go to the DevOps engineer for Command-line (CL) scripts enabling the deployment of containerized services to the edge cluster or participating clusters that have access to. Provides the minimal functional commands needed to do DevOps service life cycle on the remote side.

Policy and Trust (Slide 19)
Most difficult issue: Need sufficiently understandable and accepted trust model for resource providers to allow "external" provider to manage its service within its infrastructure. Needed for collaborative, multi-resource, coordinated infrastructure for the scale of science that we want to do in the next 10-20 years.

Deploying an "Application" (Slide 20)
Client CL and testbed deployed at core sites between Chicago, Michigan and Utah.
Natural way of looking at infrastructure.
Goal: give illusion that living in a DevOps environment (e.g., Atlas).

Summary (Slide 21)

- Goals: Reduced barriers supporting collaborative science. Ubiquitous substrate for developers. Change distributed cyberinfrastructure operational practice.
- Developing DevOps model, addressing resource providers' concerns; right tooling for developers.
- First k8s-based WAN deployments: caching servers for OSG (StashCache) to be integrated into LIGO and IceCube workflows. ATLAS: similarly deployed software (XCache). Seeking partners.

DevOps and Software Defined Networking (SDNs) and Software Defined Exchanges (SDXs) **-** *Joe Mambretti*

Introduction to iCAIR, Northwestern University (Slides 1-4)
iCAIR

- performs core research on large-scale transport networks (Layer 1 and 2) and design, implement and support national and international testbeds.
- runs production facilities: MREN- distributed Sci DMZ, SL (5600G connections and supporting over 100 private networks).
- works with the international science communities on a variety of private science networks. Drivers are large-scale science. We support virtually all large-scale science projects globally. Connected to Global Lambda Integrated Facility (GLIF), a distributed fiber optic infrastructure that is highly programmable. Can program private networks on this facility. Most of these connections are 100G.

SDN: Software is "eating" the world. (Slide 5)

- All of IT is moving to this new model with commodity resources and software as a value-add. Doing it with SDN and SDX (NSF-funded project for worldwide data intensive science).

International connections: (Slide 6)

Connects to every major R&E international connection worldwide.

Emerging U.S. SDX interoperable fabric (Slide 7)
Working with partners to develop other SDXs in the U.S.; Also some coming up in Europe and Asia.

LHC open network environment (Slide 8)
Support LHC. LHC's open network environment is used for HEP and other 6 other science communities that are not LHC communities (particle physics).

100G switching across the wide area (SC17) (Slide 9)
Working on techniques for steering large flows.

Emerging topics in advanced networking (Slide 10)
Exciting capabilities are emerging in SDN, some of which are driven by cloud provider (tenant networks).

- P4 programming language: renaissance of real programming languages. Can now get P4 switches: P4 can write code, protocols and have VLAN the next day instead of waiting 4 years (length of time for XLAN from concept to Ethereum switches). Matches accelerated pace of innovation - perfect for DevOps because you can move things quickly from virtualization to production.

Global environment for network innovations (GENI) (Slide 11)
Some of our reference projects are coming from GENI (NSF project), a distributed, highly programmable fabric with racks (each rack is a mini-cloud; considered an edge cloud).

International Software Defined Networking testbed (See Slide 12)

Chameleon (Slide 13).

- Involved in the networking pieces of this project.
- Interesting innovations in this area (e.g., Google must support 1M individual tenet networks in its environment). Google is creating tools based on OpenStack that allow you to define, implement and support tenet networkings. Doing this in Chameleon, which is based on OpenStack. Supports hundreds of individual tenet networks, as in GENI.

Bioinformatics SDXs network (Slide 14)- Working on bioinformatics SDX; SDX is a large-scale virtual switch within which you can reversibly grade virtual switches at virtual exchanges. Achieved for bioinformatics for precision medicine (2017 paper).

Working on disaggregation of the optical layer, which is normally static. Providing software overlay to make it very elastic. (Slide 15)

Starlight SDX geoscience research workflow: working with Jupyter notebooks as a steering mechanism for both science workflows and network underlayer. (Slide 16)

ESA distributed archives : Will conduct demonstration for SC18 (Slides 17-19)

- SC18 will showcase much of what's described and in the diagram of booth – provision 1200G from SL to venue and 6 x 100G from Washington D.C.

NASA Goddard Space Flight center: Holds world record for disc-to-disc over the wide area; SC16: 200G disc-to-disc over the wide area. 378 memory to memory; SC18 will do 400 G. (Slide 20)

Dynamic Distributed Data Processing (Slide 21)
Distributed large-scale applications from small activity to large scale with NRL (including using SC on 100G channels).

A Cross-Pacific SDN testbed: Big Data Express (Fermilab): specialized set of software for DTNs that understands transfers of large collections of files and large individual files. Transpacific testbed (Slide 22).

Transferring large-scale airline data E2D across WANS using DTNs (Slide 23)

100Gbps DTN optical testbed**:** useful for planning and experimenting with new techniques and protocols for 100G single flows. (Slide 24)

Designing P4 testbed as part of GENI project. Can connect anyone who is interested in using it. (Slide 25)

Open storage network (Slide 27)

- Will support science data in a distributed fabric. May be an opportunity to work with Rob Gardner (SLATE).
- Will also become part of the national research platform (part of PRP); opportunity for specialized caching and distribution services and specialized channels for optical interconnects with high performance file systems.

Software Assurance Marketplace (SWAMP) – A Continuous Assurance Platform **–** Von Welch
Team: Miron Livny (Morgridge), Bart Miller (UW) and Jim Basney (UIUC)

Broad Mission: More secured software (Slide 2)
Driving Principles: (Slide 3)

- SWAMP Framework pulls together different pieces of software assurance technology.
- Community focus: helps software assurance and development communities.
- Open source (GitHub).
- Intended to be honest facilitator and technology neutral regarding software assurance tools, makes existing tools easier and more effective to apply to projects (e.g., software development).

Software Security: It is (almost) all software! (Slide 4)
In the context of critical infrastructure, almost all infrastructure has software controlling it. Ties to ability of software to prevent kinetic effects.

A Kinetic Effect (Slide 5)
Software running shipping ports. If software is compromised, could have severe kinetic impacts.

Many Excuses (Slide 6)

- Software assurance are good, but it is not done as often as it should. Incentive issue: software security is nonvisible functionality, so not much incentive to integrate.
- Too labor intensive when integrating software assurance tools into build chain.
- Expertise required for effective usage of tools: false positives.
- Vendor lock-in: difficult to change.

Continuous Assurance**:** Assess early and often (Slides 7-9)

Looking to package available open source and commercial tools for continuous assurance (building off continuous integration paradigm; continuous tests for software strengths and lack of weaknesses in addition to build tests). Need something that integrates easily with continuous integration and build frameworks and facilitates software assurance work.

2 ways to bring the continuous assurance capabilities to the developer:

1) MIR-Swamp facility- in the cloud web-based presence; use InCommon to log in and upload software or point to github repo and do software assessments. Good way to try out SWAMP and use for software education- integrated into software development and classes.
2) SWAMP-in-a-Box (SiB)- For local deployment: open source distribution installable on-premise; no need to send software off local network.
    - Demand for on-premise by those with closed software projects; confidentiality/privacy needs.
    - Customizable (Slides 9, 13)
    - Can be integrated in local identity management system

Bring your own license: support for commercial tools.

Integrating into the SDLC (Slide 10)
Support for Integrated Development Environments and Continuous Integration/Development. Facilitates process for software developer

- E.g.: Part of it can be Integrated into the DevOps environment to provide a check before deploying as part of a continuous integration stream.

By the Numbers and Capabilities (SWAMP usage over past 4 years) (Slide 11 -12)

- Over 2 dozen open source static analysis tools and 4 commercial tools.
- Supports different platforms (software builds differently on different operating system platforms).
- Supports over 10 programming languages.
- Over 10k curated software packages with identified software weaknesses: good for checking if the software assurance tool works in SWAMP.

Road Ahead (Slide 17)

- Evolve software security training and education activities.
- Many pending requests for infrastructure support for SiB installations, different functionalities and languages (e.g., support for Microsoft windows).

Contact URLs: https://continuousassurance.org/ (Slide 18)

Discussion
Are certain types of scientific projects are more suitable to DevOps?

- 2 particle physics communities. But many science domains stand to benefit.
  Those collaborations that are building workflows to work over distributed facilities or resource providers all stand to benefit from SLATE.
- If large-scale, complex and distributed, need to go to DevOps as nothing off the shelf is suitable.
- Another consideration: Long-lived.

Reproducibility in science and extending science analyses to new models:
CERN is employing continuous integration/DevOps as relates to reproducing in-stage data analysis. Early on, but have developed interesting tools.

- In one project, studied data and software preservation for open science: discussed data preservation and managed custodial usability of data. Going forward, includes software and software environment. Now, shift going forward – multiple projects using technique.

Cybersecurity and reproducibility: DevOps is good for cybersecurity; allows you to patch and perform services quickly.

- Challenge: Reproducibility perspective where software is more fluid; infrastructure could change in the midst of a science run. DevOps requires much monitoring/debugging information in order to quickly address the deployment of new software, so fixing it "on the fly" – could help with reproducibility. DevOps requires awareness of in order to sustain frequent refreshing.

SWAMP does not go up meta layers.

- E.g., individual software package security assurance, image security assurance and then deployment or chart level assurance in subsequent layers.
- What is image security? Right now it's trusted provider image security. No overall definition of "secure image". 3 parties involved: infrastructure operator, user and internet users writ large. Need crisp definition in order to answer question.

SWAMP is focused on creating platform, not assessment tools. Conduct automation and management of process. If there are any tools to evaluate the quality of containers or other packages, we may be able to integrate them into the platform.

- External group can potentially use SWAMP's curation platform. (E.g., performing assessment or running process outside) Go to SWAMP - act of assessment and meta data going along with assurance can be curated and recorded.
- SWAMP's biggest challenge with static assessment tools: large number of false positives and most programmers do not know what to do with generated reports. Lack resources to help programmers to do something with report. How do we go from a tool to something of value to a deployer or developer of software?

MAGIC Co-Chair Nomination

MAGIC Co-chair vacancy was created when Rajiv Ramnath rotated out of NSF. Email was sent out a week ago requesting nominations of MAGIC co-chair from any interested federal agency candidates. 1 nomination received for Dr. Vipin Chaudhary (NSF); his biography was circulated and presented at the meeting.

NITRD Co-chair Nomination Process:

Team is to discuss candidate(s) and forward to NITRD Subcommittee Co-chairs, Kamie Roberts (NCO) and Erwin Gianchandani (NSF), to make the final decision.

Discussion: Candidate is well-qualified and should be forwarded to the NITRD Subcommittee co-chairs. Joyce Lee will inform the co-chairs and inform MAGIC of their decision.

CY19 Tasking Discussion

Draft Potential tasking will be presented at LSN's Annual Planning Meeting later this month for inclusion in the Supplement to the President's Budget. The draft tasking will be circulated via email to the Team.

Draft tasking:

- Task 1: Data life cycle series (4-6 months); includes provenance and potentially tools and services involved in managing complex scientific infrastructure and tools.
- Task 2: Single topic issues (e.g. identity management, edge services).

- Task 3: Engaging community through roundtables (e.g., academic community).

Potential Workshop on Data Life Cycle
Develop and provide recommendations on addressing data life cycle issues in the context of distributed computing environments.
Structure: Relying on MAGIC members to present. Organized as a series with follow-up Q&A discussion.

Roundtable
November 11 – 16, Supercomputing18, Dallas, TX
November 12, SC18 Data center automation workshop will cover data center automation technologies, from server control to provisioning systems.

Next meeting
November 14 (SC18): Kay Bailey Hutchinson Convention Center, Dallas, 1:30-3:30pm CT/2:30-3:30 ET, Room D175. Theme: Global Identity management.