ISSA JOURNAL

April 2021
Volume 19 Issue 4

Cyber Warfare: An Era of Nation-State Actors and Global Corporate Espionage

Table of Contents

DEVELOPING AND CONNECTING CYBERSECURITY LEADERS GLOBALLY

Feature

Also in this Issue

The ISSA Journal (1949-0550) is published monthly by
Information Systems Security Association
1964 Gallows Road, Suite 210, Vienna, VA 22182
+1 (866) 349-5818 (local/international)

Greetings ISSA Members

Candy Alexander, International President

Spring is here!

April is finally here and with it brings spring! Spring is a special time of year, the beginning of new life and rejuvenation. The global COVID-19 pandemic has brought us many challenges over the past year. I am sure you will agree that we are ready for a bright new season with a bit more "back to normal" activities.

I am optimistic that this spring will bring the rejuvenation of many things, including increased member participation and involvement in ISSA International projects, including the refresh of the Cyber Security Career Lifecycle ® through a review/update of the knowledge, skills, and abilities (KSA) matrix. For those of you who recall this effort, 76 volunteers worked to define KSAs for the 5 phases of the lifecycle. And it is now time to call on volunteers to help review and update them again. The information will help guide professionals through a career path, align to other sources of KSA matrices, and help identify what content and programming ISSA should develop to meet our members' collective needs.

ISSA International has plans to increase member participation opportunities, including "calls to action" to join committees and Special Interest Groups (SIGS). By participating in the committees and SIGs, we hope to share the knowledge and information with our global community.

Think of it as crowd-sourcing knowledge and pushing it back to the membership for consumption.

The ISSA International Support Team and the International Board of Directors are looking at additional ways to assist our Chapters in delivering knowledge-content back to the local membership community – through new and innovative delivery methods.

Many Chapters are finding it difficult to provide for their members through traditional webinar programs – and that is where ISSA International can help. We are researching ways to provide the infrastructure and framework for a knowledge delivery mechanism that will benefit the community on both a local and global level.

I hope that all ISSA members are as excited as I am for the rejuvenation the spring season brings. But I must also remind you that the ISSA is nothing without the active participation of its members. Whether at the local Chapter level or the International level, ISSA's success depends on you. Your support in joining the committees or SIGs is crucial, and the value you will receive is limitless. The knowledge and experience you will gain are priceless.

As always, if you have any questions or constructive suggestions, please feel free to contact me, any International Board member, or our executive director, Marc Thompson.

Sincerely,

Candy Alexander, CISSP CISM
ISSA International President
Candy.Alexander@ISSA.org

DEVELOPING AND CONNECTING CYBERSECURITY LEADERS GLOBALLY

ISSA
Information Systems Security Association

The Information Systems Security Association, Inc. (ISSA)® is a not-for-profit, international organization of information security professionals and practitioners. It provides educational forums, publications and peer inte raction opportunities that enhance the knowledge, skill and professional growth of its members.

With active participation from individuals and chapters all over the world, the ISSA is the largest international, not-for-profit association specifically for security professionals. Members include practitioners at all levels of the security field in a broad range of industries, such as communications, education, healthcare, manufacturing, financial, and government.

The ISSA international board consists of some of the most influential people in the security industry. With an international communications network developed throughout the industry, the ISSA is focused on maintaining its position as the preeminent trusted global information security community.

The primary goal of the ISSA is to promote management practices that will ensure the confidentiality, integrity and availability of information resources. The ISSA facilitates interaction and education to create a more successful environment for global information systems security and for the professionals involved.

ISSA JOURNAL

Now Indexed with EBSCO

Editor: Jack Freund, PhD
editor@issa.org

Advertising: vendor@issa.org

Services Directory

Website
webmaster@issa.org

Chapter Relations
chapter@issa.org

Member Relations
memberservices@issa.org

Executive Director
execdir@issa.org

Advertising and Sponsorships
vendor@issa.org

The Next Generation Security Professional (NGSP)

By Randy V. Sabett – ISSA Distinguished Fellow, Northern Virginia Chapter

Unlike CISA, the unfortunately overloaded acronym in the security world (standing for both the Cybersecurity Information Sharing Act and the Cybersecurity and Infrastructure Security Agency), NGSP won't have the same problem. The only thing I found on a quick search was the National Glycohemoglobin Standardization Program. Somehow, I don't think there will be any confusion of that with the new group of cybersecurity professionals on whom we will rely to protect our networks.

So, to express what I think about the next generation security professionals, I will start off with a song adapted to the tune of *Kodachrome* (which at least some of today's NGSPs probably won't even remember):

When I think back on all the cyber
I learned in high school,
It's a wonder I can think at all
And though my legal education
Hasn't hurt me none
I can read the writing on the wall

What does this have to do with the NGSPs? Well, when I was a young security professional (back when years began with 198x), I and my co-workers shared a number of common characteristics: we lived, ate, and breathed security… and played a lot of softball. My point in saying that after my Paul Simon interlude is that the information security profession continues to undergo change. Gone are the days when security professionals could just write code and play softball (though we actually did a lot more than that!) Many of today's NGSPs will need to be part programmer, part architect, part lawyer, part auditor, part salesperson, part translator, and part privacy expert.

Why the multiple hats, you might ask? Well, before people paid attention to security, the professionals in this space could pretty much dictate the what, where, why, and how of the security solutions that they were producing and implementing. This meant that very few other people could understand what the security people were doing. At least some folks noticed the danger of this and innovated in ways that caused security professionals to "learn new tricks" and broaden their view of the world. As just one example, I taught a course in the engineering school at George Washington University called Information Policy. One of the founders of the course explained to me that he wanted the course to be a combination of constitutional law, intellectual property law, privacy law, and cyber law.

The goal of my course was to expose up-and-coming engineers and programmers to numerous non-technical policy issues that clearly impact today's dialog. This should also be the goal for NGSPs today. The ability to anticipate the legal and policy impacts of tomorrow that are based on design and security decisions made today is a valuable quality that will increase security. A security professional today must be able to excel at a wide variety of things, including security awareness and promotion, proactive security activities, balancing of security and privacy, implementation of security by design principles and requirements, investigation of security incidents, handling of any security deviations, and overall enforcement of policy.

Having said all this, I'm also a firm believer in basic blocking and tackling. NGSPs must also be able to perform at the top of their game, in whatever security specialty they are working. Such specialization can make the landscape even more complicated for the NGSP, since the other dynamics (legal, policy, privacy, etc.) create even more nuanced issues.

Being a successful NGSP, including the ability to integrate the wide variety of other influences discussed above, can bring incredible professional satisfaction. As I know many of you have heard me say, for me, combining a first career as a crypto engineer with a legal second career has been incredibly rewarding and I would encourage any of you who are interested in doing something similar to explore it further. The legal community needs more folks who understand the nuances of security.

About the Author

Randy V. Sabett, J.D., CISSP, is an attorney with Cooley LLP (www.cooley.com/rsabett), a member of the advisory boards of the Georgetown Cybersecurity Law Institute and the RSA Selection Committee, and is the former Senior VP of ISSA NOVA. He has completed FBI Citizen Academy training in 2017, was a member of the Commission on Cybersecurity for the 44th Presidency, was named ISSA Professional of the Year for 2013 and an ISSA Distinguished Fellow in 2018, and can be reached at rsabett@cooley.com. The views expressed herein are those of the author and his colleagues that contributed to this article, and do not necessarily reflect the positions of any current or former clients of Cooley or Mr. Sabett.

Celebrating ISSA Women

By Dr. Curtis Campbell, ISSA Fellow, Chattanooga Chapter

This article highlights ISSA women for their resilient leadership and contributions to the field.

March marked Women's History Month. It also marked the one-year anniversary of the COVID-19 pandemic. What do the two events have in common? Both illustrate the resilience, strength, determination, and perseverance of our nations and leaders to rise-up despite some very serious challenges.

So, this month's column is dedicated to celebrating ISSA women for their ongoing contributions within ISSA, working through cybersecurity challenges, and protecting our nation's data, especially during the pandemic. We raise our hats to the great momentum our ISSA women have achieved through their chapters around the globe.

ISSA women are engaged. No barriers hold them back. ISSA is strongly represented in thirty-six countries. ISSA women are represented in 138 active ISSA chapters. At the chapter level, three of the thirty-six countries have elected women chapter presidents. In the United States, nineteen women have been elected to serve as chapter presidents.

To understand ISSA's reach, combined ISSA chapter membership totals 7,000 ISSA members in good standing. Last year, over 17,000 combined ISSA members and their invited guests attended ISSA's webinars, events, and conferences. Talk about spreading the word!

Hats Off to Women ISSA Leaders

Over the past thirty-six years, ISSA has become known for its reputation in cybersecurity leadership and expertise. ISSA events and conferences are known for providing a wealth of knowledge, networking, and educational opportunities.

Credit is due to one woman who had an early vision for growth and momentum in our field. Sandra Lambert, founder of ISSA, saw the need for an organization to educate and become the trusted voice of the cybersecurity industry. Known for her strong leadership and industry expertise, Sandra's influence and achievements span over three decades. Thanks to her vision, so have ISSA's.

Sandra Lambert was also elected as the first international ISSA board president to chart ISSA's course. Over the years, ISSA has elected and re-elected 17 international ISSA presidents to date. Seven women ISSA international presidents have been elected including our current ISSA international board president, Candy Alexander. A 30-year veteran cybersecurity leader and CISO, Candy Alexander is currently serving her second consecutive 3-year term, a testament to her industry knowledge and strong leadership.

In addition, two ISSA women serve on our current international board as vice president and treasurer/CFO, serving multi-year terms. In 2021, three ISSA women (out of eight board members) were elected as international directors on the board serving multi-year terms. These women are resilient, proven leaders.

Early beginnings

In 1984, when ISSA began, the tech world was on the verge of propelling us forward to a digital world. In 36 short years, we have seen technology transformation at lightning speed. In 1984, Dell computers launched, and the Macintosh computer aired during the Super Bowl. PCs ran DOS. In the '90s, Amazon, Yahoo! and Mosaic Communications (later Netscape) were just beginning. Computers were common but featured floppy disks and dial up. By the mid 90's, most people used Windows, and the World Wide Web was up running. In 2014, the Web celebrated its 25th birthday. Today's technology development continues in IoT, AI, Autonomous driving vehicles, and Virtual Reality to 5G, 3D Printing, Drones, Biometrics, and Quantum Computing. In the past year, we have conducted much of day-to-day business by social distancing with Zoom, so technology and cybersecurity continue to roll with the punches.

During these formative years, ISSA focused on developing a successful global environment for our security professionals, with chapters forming all over North and Latin America, Europe, and Asia. Its international vision worked: to provide members a way to connect and collaborate, expand peer networks, enhance professional stature, and achieve career goals.

Still Work to Do

Today, ISSA continues to provide value to members through programs that serve to enhance and educate us on cybersecurity challenges in the field. It is an inclusive environment for women and a great venue for narrowing the gender gap in cybersecurity. Yet, ISSA female chapter presidents represent only 13.7% of chapter leadership. While this is a good percentage, there is still some work to do. More women are needed and what better time than the current as we begin to see our way clear of a desolate and dangerous past pandemic year.

continued on page 17

Avoiding Cost Disease

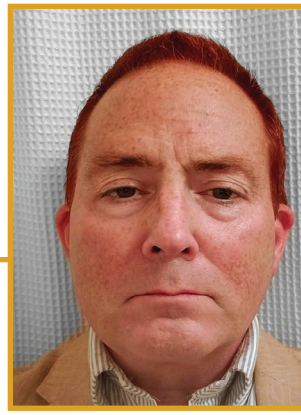

By Luther Martin – ISSA Member, Silicon Valley Chapter

Things which don't get more efficient with better technology also tend to get more expensive over time. Economists William Baumol and William Bowen noted this pattern in the 1960s. They called it "cost disease." The explanation of cost disease goes roughly like this. Suppose Alice and Bob work in different industries, and both make $50 per hour. Now suppose that better technology at Alice's company increases her productivity and thus her wage to $60 per hour. Bob might want to go work for Alice's company for the higher wage, so to keep him, his company will also need to increase his wage to $60 per hour. But because this wage wasn't justified by higher productivity, Bob's company will need to raise prices to fund the higher wage while Alice's won't.

Cost disease can explain a lot of the increases in the price of healthcare over the past few decades. Even with better technology, a doctor can only treat a limited number of patients per day. This makes it hard to increase the productivity of the healthcare industry by using technology, so the cost of healthcare tends to rise faster than the overall rate of inflation. This increase isn't caused by big bonuses paid to executives or corporate greed, it's just what economics tells us to expect.

Cost disease is also responsible for at least part of the dramatic increase in the cost of college over the past few decades. Professors can only teach a limited number of students at a time, which limits the ability to make them more productive using technology. All of the research that I have seen about this was done before on-line classes were popular, and it will be interesting to see what future research shows. I suspect that on-line alternatives will not necessarily make education more efficient. Every professor that I've talked to about it has used on-line material to supplement their face-to-face interaction instead of replacing it. Purely on-line alternatives may get very different results.

So to keep security affordable we need to make it more efficient. This has happened, but it needs to continue into the future.

Back in 1999, I did what I believe were the first commercial code reviews. If you did them before 1999, let me know. My fraternity chapter claimed to be the first fraternity house with a swimming pool. It may or may not have been true, but since nobody ever told us that it wasn't true we kept saying it. Similarly, I'll keep saying that I did the first commercial code reviews until someone tells me that I'm wrong.

In 1999 there were no tools available that checked code for application security vulnerabilities (buffer overflows, SQL injection, etc.). Without any tools, we did this by hand, using the grep command to look for vulnerable functions that might have been carelessly used and then looking at the code to see if it was vulnerable or not. I hope that nobody is doing it this way today. There are lots of tools available now that can do a much better job than four people in a windowless basement room can. The results are better and are much quicker to get. That's the sort of increased efficiency that prevents cost disease.

Security Information and Event Management (SIEM) technology is another example of a good use of technology. Back when I was using grep to look for careless uses of strcpy(), security administrators were commonly looking through event logs using similar technology to find patterns that might suggest that they had been hacked. This was hard, expensive, and very prone to error. Today, however, SIEM products do a great job of this. These products are doing a much better job much more efficiently than a person ever could. That's yet another example of the improved efficiency that will prevent cost disease.

So the security industry has come up with solutions that are moving us in the right direction, but we need to keep that happening.

Enterprise key management too often requires the involvement of administrators and these administrators make lots of mistakes. According to David Smith's Reliability, Maintainability and Risk, we can expect an error rate of about 1 percent for typical security administrator tasks (Smith's "routine with care needed" tasks). That may not sound too bad. After all, a grade of 99 percent will get you an "A" in absolutely any class. But if any one of those 1 percent of errors causes an exploitable vulnerability, then it will have left a door open for hackers. Key management is hard and expensive, so it's a good candidate for more automation in the future. And doing this will help us prevent the cost disease that might otherwise plague the industry.

About the Author

Luther Martin has survived over 30 years in the information security industry, during which time he has probably been responsible for most of the failed attempts at humor in the ISSA Journal. You can reach him at lwmarti@gmail.com.

Serene Surveillance

By Karen Martin – ISSA Member, Silicon Valley Chapter

La Serenissima, the serene city of Venice, Italy has suffered greatly this year as tourism has dropped. But for the rest of the last 20 years, it has been anything but serene. Starting in the early 2000s, it was swamped with day-trippers for 11 months out of the year. In 2015, 30 million people visited the 3 square miles of the historic center of Venice during the tourist season, an average of nearly 90,000 people per day.

The city just introduced a new weapon to manage the crowds: The Smart Control Room. The goal is understandable, but the method involves surveillance that might strike citizens and tourists as creepy. Venice is not alone in looking to surveillance for solutions to common civic problems like congestion, traffic accidents, and noise control. With the advent of connected devices like CCTV cameras, noise sensors, connected vehicles, and mobile phones, it is much easier for cities to gather data that may help address some of these issues. We should all be pressuring them, however, to make sure that these solutions are both effective at improving civic life and that they do it with the minimal possible impact on privacy, for both citizens and visitors. Venice may not be getting that balance right.

Over-tourism is certainly a huge issue in Venice. During the 11-month tourist season, pre-COVID, the islands making up the historic center of Venice would get nearly 100,000 tourists a day on average. They were mostly day-trippers, arriving on cruise ships or via the bridge from the mainland, all intent on getting from the port or the train station to the Rialto Bridge and St. Mark's Square along the same narrow walkways. Day-trippers provide little value to Venice. They may or may not purchase souvenirs or food from the vendors lining their route, but they definitely generate waste and noise, and cause congestion on the main waterways and walking routes.

In the past, tourism data was focused on overnight stays. But as day-trippers have become more common, Venice has struggled to collect accurate visitor counts and control traffic flows. The new Smart Control Room monitors the flow of traffic on the canals, identifies the types of vessels and their direction and rate of travel, monitors whether the public vaporetto water buses are running on schedule, and controls traffic signals as needed to reduce congestion on the waterways.

It also receives cell phone location data through TIM, Italy's most widely used telephone service provider, to track individual cell phones in Venice, the country of registration and, in the case of Italian-registered smartphones, the region in which the phone is most active. CCTV cameras are also available to estimate foot traffic and pedestrian speed. Real-time information about foot traffic is probably useful, as the city could use turnstiles or barriers to route visitors over different routes to the main sights, while allowing locals to move more freely.

But how does knowing the home location of all the smartphones that visit Venice prevent over-tourism? There might be a legitimate case for differentiating between smartphones that are present on most working days, which likely belong to the dwindling tribe of Venetian residents or the workers who commute daily from the mainland, and those present for short periods of time. But does it really matter where the visiting phones came from and do you need to be able to track them individually?

The Netherlands city of Eindhoven found a more privacy-centric solution to deal with noise, crowding, and violence in a popular nightlife area. It used audio sensors to report the level and direction of noise, without recording conversation and video sensors to count and report the number of people passing without storing the video. Although cellphone registration locations were tracked, the data was aggregated to the municipality level to protect privacy. If governments are determined to use censor data and track connected devices, we should at least insist they do it in the least invasive way.

But even so, high tech solutions are not always the best or cheapest solution. Venice spent about $3.5 million on its Smart Control Room, which is going to manage traffic, not limit it. They also plan to start charging a variable entrance fee for day-trippers, which will increase on crowded days. That could actually reduce congestion, which sounds like a better outcome. It will certainly cost the residents less.

About the Author

Karen Martin is a San Jose based information security engineer. She may be reached at kjlmartin@gmail.com.

The Remote Workforce Will Lead to More Ransomware Incidents in 2021

By Rusty Carter, CPO of LogRhythm

As companies adapt their business to a remote workforce amid the COVID-19 pandemic, an even larger attack surface for cybercriminals is being created. Over the past two years, and accelerating the last few months, ransomware attacks have become one of the most common threat vectors. . According to PwC's Cyber Threats Report, ransomware was the most significant cybersecurity threat of 2020 with incidents more than quadrupling over the second half of the year. Organizations ranging from healthcare facilities to education and financial institutions have found themselves more vulnerable than ever to these types of attacks.

Despite cybercriminal groups pledging to stop attacking healthcare and medical organizations, the Cybersecurity and Infrastructure Security Agency (CISA), Federal Bureau of Investigation (FBI), and Department of Health and Human Services (HHS) recently released a joint advisory after witnessing an increase of cybercrime threat to U.S. hospitals and healthcare providers.

Now is the time for enterprises to ramp up current detection and prevention strategies, particularly against ransomware, as this form of attack will continue to wreak havoc this year.

The perfect storm

The correlation between remote work and an increase in ransomware is impossible to ignore. As a bulk of the workforce switched to remote work at the beginning of 2020, the FBI's Cyber Division reported that cyberattacks had increased a startling 400% at the onset of the pandemic to as many as 4,000 incidents a day. With nearly half of organizations in HR, legal, compliance, finance and real estate industries planning to allow full-time remote work for employees going forward, this trend will undoubtedly continue into 2021 and well past the end of the pandemic.

As we've seen in recent months, cybercriminals are ramping up attack efforts as remote work remains in effect and vaccine rollouts continue to increase. These larger attack surfaces combined with new avenues for monetization has only fueled their motivation. For example, Baltimore County Public Schools experienced a ransomware attack that infected its systems and forced it to shut down most of its networks. As a result, the school district's offices, email system, grading system and website were all impacted, and the BCPS was consequently forced to suspend all virtual learning, disrupting student learning for multiple days.

Recovery efforts from a ransomware attack can also be complex and lengthy, especially if organizations experience significant data loss due to inadequate backup procedures. Hackers know this impact is detrimental and that a company may easily be influenced into paying the ransom so the business can get back up and running. However, the decision to pay the ransom should not be made lightly seeing as it could result in additional fines from the U.S. Federal Government.

Unfortunately, the growing number and distributed nature of devices and data cannot be slowed down, as digital consumption and expansion exponentially increase. Additionally, public cloud services are forecasted to expand nearly 20% in 2021 and control continues to become more difficult, while systems leveraging things like microservices have exponentially more interfaces that become the target of attack. Together, these shifts have forced companies to face the growing challenge of adequately defending themselves against a rapidly changing security landscape and flurryof cyberattacks. To combat this, companies must increase their focus on detection and response.

How companies can prepare

Detection and response early in the cyberattack lifecycle is key to protecting the company from a large-scale impact. As there will likely be an uptick in attacks over the coming months, it's critical that enterprises can rapidly respond to a gap or vulnerability. The below security controls will enable organizations to detect an intrusion in real-time and allow the security team to spring into action and respond.

Visibility: Ransomware is not a one and done kind of attack. In reality, ransomware attacks can actually happen several times to the same company, especially if the right precautions are not taken after the initial attack. No matter where an organization stores its data, real-time monitoring and clear visibility are crucial for rapidly detecting and neutralizing security threats. A comprehensive view of the endpoint and server activities allows the security team to detect and monitor suspicious activity. With this type of constant visibility, companies know if they are secure or if there is a vulnerability.

Identity: Companies also need to be able to respond quickly to anomalous activity. User and entity behavior analytics can help monitor known threats and behavioral changes in user data, providing critical visibility to uncover user-based threats that might otherwise go undetected. This includes insider threats, compromised accounts, or privilege misuse.

Continued on page 17

News from the Foundation

We are delighted to welcome our new Board Director, Shelley Wark-Martyn, who began her term on February 8, 2021.

Shelley Wark-Martyn has been instrumental in connecting companies and individuals to the training needs and programs SANS provides through intensive, immersion training. In accepting the invitation to join the Foundation board Shelley said "Excited to volunteer for ISSAEF, creating more opportunities for scholarships and training. Looking forward to continuing the work in creating a strong training foundation for the future."

Shelley steps in to fill the position vacated by Randy Sanovic who retired last December. The Board regrets his departure and extends their appreciation for his work during his tenure to further the Foundation's mission and goals.

CORPORATE SUPPORT

Thank you to The ISSA Colorado Springs Chapter for making a generous contribution in 2020 to the ISSAEF scholarship fund.

Cloud Security Alliance the world's leading organization dedicated to defining standards, certifications, and best practices to help ensure a secure cloud computing environment has provided ISSAEF with grants totaling up to $2,000 to be awarded in 2021 to support and strengthen the cybersecurity profession. Jim Reavis, CSA CEO, stated, "We are grateful for ISSAEF's efforts to further the development of professionals within our industry." Check out the eligibility criteria and application deadline at issaef.org.

Once again, SANS has entrusted ISSAEF with offering access to one of their exceptional SANS courses. This new application process for the E. Eugene Schultz, Jr. Memorial Training Scholarship broadens that availability across North America. Individuals can target specific courses of interest to accelerate their cybersecurity training. Check out the eligibility criteria and application deadline at issaef.org.

2021 CYBERSECURITY SCHOLARSHIPS AND GRANTS

The application period for the 2021 ISSA Education Foundation Scholarships that opened on February 1st and will remain open through the dates listed below. The application form is available at https://issaef.org/scholarships/ with the following scholarships/grants available:

Application Deadlines June 15, 2021

- Howard Schmidt Memorial Scholarship for undergraduates - $3,500
- E. Eugene Schultz, Jr. Memorial Training Scholarship for graduates - $3,500
- Shon Harris WIS Memorial Scholarship for women in security - $2,000
- Alamo ISSA Chapter Scholarships - George "Chip" Meadows Memorial Scholarship ($1,500), two $1,000 and one $500 scholarship, all for local universities

Application Deadlines April 30, 2021

- Cloud Security Alliance (CSA) Professional Development Grant - $2,000
- SANS E. Eugene Schultz, Jr., Memorial Training Scholarship – One training

OPPORTUNITY TO JOIN THE ISSAEF BOARD

The Foundation is seeking a new Director to assist with Special Fundraising programs. If you have experience serving on Not-for-Profit boards with activities such as silent auctions, conference-related drawings, or similar fundraising opportunities and are interested in joining the Foundation board, please write to president@issaef.org.

SEEKING VOLUNTEERS FOR SCHOLARSHIP AND GRANT COMMITTEES

The Foundation is looking for Cybersecurity professionals to volunteer on the upcoming 2021 Scholarship Review Committee and Professional Grant Review Committee. Are you interested in paying it forward? We need your assistance in evaluating grant awardees in May 2021 and scholarship applicants in June 2021 for the current year award cycle. Please contact volunteer@issaef.org and let us know your interest and background.

We are also seeking volunteers to participate in short term projects, scholarship publicity, fundraising, and governance of the Foundation. Those interested in joining a truly dedicated and enthusiastic group, please send an email with your background to volunteer@issaef.org

SUPPORT US WHILE SHOPPING

Help spread the word about these great opportunities to your friends and family at no cost to you – just use Amazon Smile while shopping online and automatically, with absolutely no cost to shoppers a 0.5% of eligible purchases will be donated by Amazon to our scholarship fund! It's simple: start the purchase on https://smile.amazon.com, select "ISSA Education and Research Foundation Inc." (needs to be done only the first time), and shop as usual. Do not forget to tell your family/friends to do the same.

Like us on Facebook and LinkedIn

Support us with amazonsmile You shop. Amazon gives.

ISSA Community Corner

Awards Program Open. Now accepting nominations

Every year ISSA International recognizes excellence in information security professionals, the companies they work for, and the chapters to which they belong.

Find out more at

https://www.issa.org/issa-international-awards-2021/

Fellows/Distinguished Fellows

The ISSA Fellows Program honors established cyber professionals with demonstrated success and contributions to the industry. These individuals have dedicated years towards the innovation and progression within the cyber realm.

Find out more at https://www.issa.org/fellows-program/

Our 5th annual ESG/ISSA Research Survey is open now

We are surveying cybersecurity professionals like yourselves to better understand the cyber-landscape and how it has affected our profession, our careers, and our organization's security posture.

- Cybersecurity careers
- Skills development
- Cybersecurity organizational considerations
- Security incidents and vulnerabilities
- The cybersecurity skills shortage
- Cybersecurity activities

Join your cyber security colleagues from around the world to ensure your perspectives are earmarked for further development and analysis.

Take an active role in the future of the cyber security industry and the InfoSec careers!

Take Survey Here https://survey.az1.qualtrics.com/jfe/form/SV_231qBzOnmFllUzA

Chapter Leader Meetings Schedule

- March 26, 2021 - 1:00 PM Eastern Time
- April 20, 2021 - 1:00 PM Eastern Time
- May 28, 2021 - 1:00 PM Eastern Time
- June 25, 2021 - 1:00 PM Eastern Time

Find out more at https://www.members.issa.org/page/ChapterLeadersSummit

Upcoming Events and Conferences

Cyber Executive Forum – Virtual Summit

May 14, 2021 is the next Cyber Executive Virtual Summit.

We welcome guests (ISSA general members and non-members interested in the Cyber Executive Forum and Cyber Executive Membership) to put in an application to join these sessions.

All ISSA Cyber Executive Members are invited to attend and invite guests.

Find out more at https://www.issa.org/event/may-virtual-cyber-executive-forum-2021/

ISSA Members receive discounts to a variety of industry events and conferences such as:

- RSA Virtual Conference
- InfoSec World Conference

To learn more visit: https://www.members.issa.org/general/custom.asp?page=SpecialOffers

ISSA Journal 2021 Calendar

JUNE
The Infosec Toolbox: Basics to the Bleeding Edge
Editorial Deadline 5/1/2021

JULY
Security vs Privacy Tug of War
Editorial Deadline 6/1/2021

AUGUST
Disruptive Technologies
Editorial Deadline 7/1/2021

SEPTEMBER
Shifting Security Paradigms in the Cloud
Editorial Deadline 8/1/2021

OCTOBER
The Business Side of Security and Risk Management
Editorial Deadline 9/1/2021

NOVEMBER
Big Data/Machine Learning/Adaptive Systems
Editorial Deadline 10/1/2021

DECEMBER
Looking toward the Future of Infosec
Editorial Deadline 11/1/2021

For theme descriptions, visit
www.members.issa.org/page/journal-editorial-calendar

EDITOR@ISSA.ORG • WWW.ISSA.ORG

Cyber Warfare: An Era of Nation-State Actors and Global Corporate Espionage

By Joab Kose – Incident Response & Security Researcher

The last decade witnessed the highest rise in state-sponsored Cybercriminal activities, with the most recent Russian-linked SolarWinds breach that impacted most of the US governmental and non-governmental organizations, confirming how dangerous these threats can be.

Executive Summary

The last decade witnessed the highest rise in state-sponsored Cybercriminal activities, with the most recent Russian-linked SolarWinds breach that impacted most of the US governmental and non-governmental organizations, confirming how dangerous these threats can be. These types of attacks are performed by Advanced Persistent Threats (APT) groups, with specific missions on their target victims. APTs are hacking groups that are sponsored and funded by governments. The groups are well-organized, highly skilled, experienced, and determined. They have strategic modes of operations. It is publicly known that countries like Russia, China, Iran, , the USA, Israel, and North Korea, own hacking groups that are trained and well-resourced to carry out specific missions in the interest of their countries. This article takes an in-depth look at the cyber threats from the state-sponsored cybercriminal groups and cyber-espionage activities they are involved in. I must state that this article reflects my own opinion, based on studies, research, and other articles that have been published, and not the publisher or the association's opinion.

Keywords

Advanced Persistent Threats (APT), Cyber-Espionage, State-Sponsored Attack, China, Russia, Iran, North-Korea.

Introduction

In 2019, the world was shocked by the revelation about the C919 airplane which was manufactured by a Chinese state-owned aerospace company known as Comac. It turned out that the airplane was a product of a Chinese global hacking operation leading to the illegal acquisition of intellectual properties (IPs) of different parts of the plane, from several foreign companies like Ametek, Honeywell, Safran, Capstone Turbine, and GE, among others, between the year 2010 and 2015. It is reported that this was a multi-year well-coordinated hacking-campaign sponsored by the Chinese Government, in its attempt to bridge the technological-gap in China's aviation industry. [7] While this alone might sound scarey, China is not the only threat actor in this kind of operation. Russia, Iran, and North Korea have proven their nation-states' cyber-capabilities through offensive operations like discovering secrets, stealing corporate data (intellectual property), corrupting individuals through political disinformation, spying on specific targets, disrupting operations, and destroying critical infrastructures of other nations. [6] However, China has been the main player in cyber-espionage, which is the focus of this article.

With the realization of the opportunities that cyberspace is offering, governments across the world are building on Cyber-forces, which are tasked with accomplishing their specific goals and agendas against other states. [4] Nation-state actors are well-trained, resourced, and equipped to disrupt, steal, and interfere with other nations' economies, governance, and military capabilities. APTs differ from other criminal hackers in several ways: they are normally not interested in personal gain, and they engage in long-term cyber operations that could go undetected for several months or years. Protecting from APTs is challenging because their attacks are directed at several security layers of their victims, and they are super stealthy in their modes of operation. They deploy some of the most sophisticated techniques, tactics, and procedures that can go undetected by the security layers in place.

What is motivating Nation-state actors?

There has been an ongoing struggle for superiority, dominance, and relevance among the developed and developing countries globally. Superiority and influence are determined by certain aspects, like military strength, economic success, and the resources that a country possesses. For years, there have been

restrictions, borderlines, and boundaries to dictate what one country could do to another. However, the advancement in technology has broken these barriers, and Cyberspace has offered the opportunity to those with the capability and willingness to utilize and take advantage.

To gain economic and technological power, some countries have resorted to climbing their way up the economic ladder through illegal means and shortcuts. This involves stealing the cutting-edge technologies and innovations from other countries that have invested a lot of their resources and time in research. There are key fields of interest like healthcare, aviation, military technology, among other sectors, in which a lot of intellectual properties are being stolen. China remains a bigger player in corporate espionage, and it uses all means, including cyber-intrusions and corrupting corporate insiders to gain access. [11] This is in addition to the disinformation and other cyber-criminal campaigns being carried out against other countries. The drive behind economic espionage to gain economic power and cyberspace is providing the platform for all these to happen.

According to Cimpanu, [7] China invests a lot of resources in the illegal acquisition of intellectual properties from different companies and institutions across the globe to achieve its economic and technological goals. The Chinese nation-state actors carry out their coordinated hacking campaigns, and sometimes when they hit the dead-end and unable to accomplish their missions through cyber-intrusion; they switch to corrupting some of the trusted insiders from their target companies. This has also been witnessed in the higher education and research institutions where some countries corrupt trusted researchers and graduate students to carry out their corporate espionage missions. [11] China is accused of sending its students and researchers as visiting scholars, to International research institutions, and using them to spy and steal research work across the globe. [3] According to the China Defense Universities Tracker, [5] China has a list of universities that are linked and integrated into the Chinese Military apparatus, intelligence community, and security agencies across China. Sending students and researchers from these institutions to foreign countries as visiting scholars creates the bridge for the wider espionage campaign.

How dangerous can the threats from Nation-State actors be?

In 2007, Estonia was hit by one of the deadliest and politically motivated cyber-attacks in history, and experienced an internet blackout for several days, exposing the capability and threats of state-sponsored cyber-forces. It later came clear that the attack was a result of the political conflict between Estonia and Russia, and it is believed that Russia was responsible for the attack. Estonia, being a small country with most of its activities digitized and connected to the internet, the DDoS attack disrupted critical operations like banking, transportation, and communication systems for days. [9] Three years down the line, Iran became the victim, with a deadly malware attack: Stuxnet, which was targeted at its Nuclear facility (Natanz uranium enrichment plant). This was a well-researched and perfectly executed attack that targeted a specific component used for controlling the centrifuges. Stuxnet became the first digitized weapon used against another nation, with the impact of bringing down the entire nuclear plant. [10] According to Rosenbaum, [13] the US and Israel possibly played a role in the Iran Stuxnet attack. In retaliation to Iran shooting down the US drone in 2019, it's believed that the US responded with a cyber-attack that disabled Iran's computer systems used to control missiles and rockets' launchers. [14]

Nation-state actors pose a huge threat to their targeted victims because these APT groups deploy techniques, tactics, and procedures that have the potential of causing damaging impacts. This became very clear in 2014 when SONY got hacked by the North Korean state-sponsored hackers following the release of the movie "The Interview," in which North Korea claimed made fun of their president. Again, this was a well-executed attack that showed how governments have invested in their cyber-forces and capabilities to carry out specific missions. These instances show the impact and threats being posed by state-sponsored cyber-forces, and how far they are sometimes willing to go.

Targeted areas of interest by Nation-State actors

Nation-state actors have specific goals and areas of interest in their operations, and they invest heavily to succeed. Some of their missions are long-term and require more resources and skills to achieve. Others are short-term and instant, but still, need investment and sponsorship from their states. Each nation-state actor has a different interest and motivation, and research has shown that the major interests in foreign countries are intellectual properties (IPs), Political and governance interference, and military technology.

Intellectual Properties (IPs)

There has been a mass campaign for economic espionage from different countries, targeting specific sectors in selected countries. Many Chinese nationals were charged by the United States Department of Justice in 2014, in connection with corporate espionage against the United States corporations. Most of the charged victims were Chinese state officials working in different units from the PLA (Chinese People's Liberation Army). [2] China is well known to have perfected its art of espionage and has several APT groups tasked with stealing intellectual properties across the globe. Their main interest is cutting edge innovations in healthcare, technology, aviation, and transportation. The targets are mainly big companies across the globe that work on research and new cutting-edge technology products relating to satellite-industries, aerospace, and communication-industries. Russia is known to have built some of the powerful tools for cyber espionage. These tools include Mini-Duke, Cosmic-Duke, Onion-Duke, and Cozy-Duke, and are believed to have been built and used for cyber-espionage by a Russian Hacking group known as the DUKE. [4]

Political and Governance Interference

Foreign countries are becoming more interested in the decision-making process of other countries that have influence. This is a wider global cyber-campaign, mainly being led by Russia. [1] Through online disinformation, cyber intrusion, and corrupted insiders, the state-sponsored hacking groups are trying to interfere with politics and governance in other countries to gain global political mileage. The last decade recorded the highest political influence through online campaigns, and the revelation about Cambridge Analytica was just the tip of the iceberg of how decision making can be influenced through electoral processes. The Estonia hack, allegedly by Russia [9] in 2007 showed how political interests and conflicts can be a major precursor to massive cyberattacks, leading to losses and destruction of assets.

Military capability

A Chinese APT group by the name PUTTER PANDA has shown a lot of determination in conducting reconnaissance and intelligence gathering missions, with the United States as their target. According to CrowdStrike, [2] this group is targeting the United States Defense, Research institutions, and the technology sectors. The Defense contractors have been the main target for cyber-espionage, with British defense contractor QinetiQ having been compromised by an APT group linked to China. During this breach, the attackers were able to gain access to information about the United States' cutting-edge military drones and robotic-weapon systems. [8]

Attack techniques used by Nation-state threat actors

State-sponsored cyber-attacks are performed with greater precision. A lot of effort, research, and resources are invested before the real attack. Some of the techniques and methods used by the attackers include:

Cyber-Intrusion

Cyber intrusion is the main method used for stealing corporate intellectual properties and assets. This involves compromising the target systems and networks to gain remote access to obtain the data they need. The hacking groups have smart, skilled, and experienced personnel, with sophisticated tools that they used to compromise their targets. In most cases, they look for zero-day vulnerabilities to exploit.

Corrupting Trusted Insiders

On several occasions, the APT groups perform massive campaigns to corrupt the trusted insiders, to gain access to the organizations' information. Research has shown that most of the breaches are linked to the people who work with the companies and organizations being targeted. This type of attack is hard to detect and protect from because the people being used by the adversaries have approved access to the targets. During the Chinese global hacking operation between 2010 and 2015, that led to the acquisition of intellectual properties from different companies across the globe to build the C919 airplane, it is reported that when they could not obtain what they wanted through cyber intrusion, they could turn to bribe the people who worked in these companies [7].

Using graduate students and researchers

This is another method being used by foreign countries to steal research materials from research institutions. Research Institutions have become the focus and easy targets that admit foreign students and researchers under programs such as visiting scholars. The FBI charged Harvard's Chemistry department chairperson for having given false information that related to the Chinese talent-plan, and the PLA-Officer who was admitted at Boston University. It turned out that the PLA officer posed as a student while spying. FBI also reported that they arrested a Chinese-researcher who was stealing and smuggling biological-research vials in Boston. [11]

Disinformation through social media platforms

With the increased usage of social media platforms, the internet offers a cheaper and faster way of reaching large masses of people. Disinformation is another method used by foreign actors in their attempt to achieve their ill-intentions. The revelation about the Cambridge Analytica campaign to change peoples' views and decisions through the provision of misleading information, proved just how online-based disinformation could be used to change the course of countries' ways of life and reasoning.

Challenges from Nation-state attacks, and the way forward

Based on the trends that have been witnessed with the nation-state operations in the past years, it remains a challenge to protect organizations from state-sponsored Cyber-attacks and corporate espionage. Most of the known attacks have been noticed and detected after the breaches and damage. APTs deploy some of the most sophisticated methods in their hacking operations, making it harder and more difficult to be detected during the initial stages of the attacks, and even after gaining access. The biggest challenge is that these attacks always target different security layers of the organizations, including the exploitation and corrupting the trusted insiders with privileged access. This technique has been on the rise, especially with the corporate espionage campaigns from the nation-state threat actors targeting research institutions in other countries. A successfully executed attack from the Nation-state hackers could result in big losses and massive damages, because of the resources and time that attackers invest in their missions. This became clear in the most recently Russia-linked SolarWinds hack that targeted the software development stage and went undetected for several months after many organizations were breached. [12]

Protecting organizations' and governments' assets from nation-state cyber threats requires proactive, active, and reactive security postures, in addition to the deployment of multi-layered security strategies. For instance, this could include avoiding the usage of equipment made by vendors from the suspected nations with state-sponsored actors, investing in the human aspect of the security for the organization, through constant training and security awareness. Humans can become an easy target to be exploited. Additionally, with the studies and intelligence gathered from previous nation-state cyber-attacks, there

are security frameworks that have been developed by security experts to reduce the attack surfaces of the organizations. These frameworks are not security tools, but just layouts of how effective security should be implemented. This article only highlights three of such frameworks, and it is worth mentioning that the proper implementation of these frameworks in organizations has tremendously reduced state-sponsored cyber-attacks.

NIST Cybersecurity Framework:

NIST Cybersecurity Framework defines five functions that should be implemented to track and secure an organization's assets and infrastructure. Each function acts as an implementation phase with specific requirements and practices. It is worth noting that this is just a framework and an organization should implement it in a way that meets the business requirements of the company. NIST defines five stages: Identifying, Protecting, Detecting, Responding, and Recovering. The first three stages of this framework highlight the proactive and active security postures that be well implemented to protect organizations from attacks or detect any attempted attack. The last two stages are proactive security postures which define how to respond to a security incident.

ATT&CK Matrix:

This is a knowledge-based model that works on adversarial-tactics and techniques based on real-world observations. This framework can be utilized as a foundation to develop specific threat-models and methodologies affecting the private-sectors, governments, and cybersecurity products and service-communities. ATT&CK Matrix emulates the initial access, execution, persistence, privilege escalation, defense evasion, credential access, discovery, lateral movement, collection, command, and control (C2), exfiltration, and impact, from the attacker's perspective. A proper implementation will in understanding the attackers and their techniques.

Cyber Kill Chain:

The Cyber Kill Chain framework is based on an Intelligence-driven Defense model to identify and prevent cyber-intrusion activities. It works by identifying the activities that the attackers must go through to accomplish their objectives and breaking that chain to disrupt the attackers' workflow. For a successful attack, attackers should completely progress through the attack stages defined by this model. Understanding the attackers' attack stages and breaking the chain between these stages will make it harder for the attackers to succeed. The attack stages defined by this framework are reconnaissance, weaponization, delivery, exploitation, installation, command, and control (C2), and action on objectives.

Conclusion

The emergence of Nation-state threat actors introduced a new security challenge in cyberspace, with developed and developing countries building and hardening cyber-forces and capabilities. We have witnessed what APT groups are capable of: from corporate espionage, political and governance interference, to trying to disrupt the military capabilities of other countries. With their sophisticated modes of operations and resources, nation-state cyber-threats have been successful in most of their hacking campaigns, and organizations are still struggling and having challenges in protecting their assets from the APTs. However, there are security measures and strategies that can be implemented to reduce the attack surfaces in the organization and reduce the success rate for most of the attacks from state-sponsored attackers.

References

1. Mueller, R. (2018). United States Grand Jury Indictment. Retrieved 28 January 2020, from https://www.justice.gov/file/1080281/download
2. CrowdStrike (2020). Crowdstrike Intelligence Report. Retrieved 8 March 2020, from http://cdn0.vox-cdn.com/assets/4589853/crowdstrike-intelligence-report-putter-panda.original.pdf
3. Joske, A., & Jones, C. (2019). How China Uses Its Universities to Spy on America. Retrieved 19 March 2020, from https://nationalinterest.org/blog/buzz/how-china-uses-its-universities-spy-america-100557
4. F-Secure (2020). The Dukes 7 years of Russian cyberespionage. Retrieved 28 January 2020, from https://www.f-secure.com/documents/996508/1030745/dukes_whitepaper.pdf
5. Australian Strategic Policy Institute (2020). China Defense Universities Tracker. Retrieved 19 March 2020, from https://unitracker.aspi.org.au/
6. FireEye (2020). [Report] Double Dragon: APT41, a Dual Espionage, and Cyber Crime Operation. [online] Available at: https://content.fireeye.com/apt-41/rpt-apt41/ [Accessed 28 Jan. 2020].
7. Cimpanu, C. (2019). Building China's Comac C919 airplane involved a lot of hacking, report says | ZDNet. Retrieved 5 February 2020, from https://www.zdnet.com/article/building-chinas-comac-c919-airplane-involved-a-lot-of-hacking-report-says/
8. Schwartz, M. (2013). China Tied To 3-Year Hack Of Defense Contractor. Retrieved 5 February 2020, from https://www.darkreading.com/risk-management/china-tied-to-3-year-hack-of-defense-contractor/d/d-id/1109795
9. Ottis, R. (2007). Analysis of the 2007 Cyber Attacks Against Estonia from the Information Warfare Perspective. Retrieved 2 March 2020, from https://ccdcoe.org/uploads/2018/10/Ottis2008_AnalysisOf2007FromTheInformationWarfarePerspective.pdf
10. Zetter, K. (2014). An Unprecedented Look at Stuxnet, the World's First Digital Weapon. Retrieved 2 March 2020, from https://www.wired.com/2014/11/countdown-to-zero-day-stuxnet/
11. Wray, C. (2020). Responding Effectively to the Chinese Economic Espionage Threat. Federal Bureau of Investigation. Retrieved 9 March 2020, from https://www.fbi.gov/news/speeches/responding-effectively-to-the-chinese-economic-espionage-threat
12. Microsoft Security (2021). Analyzing Solorigate, the compromised DLL file that started a sophisticated cyberattack, and how Microsoft Defender helps protect customers - Microsoft Security. Available at: <https://www.microsoft.com/security/blog/2020/12/18/analyzing-solorigate-the-compromised-dll-file-that-started-a-sophisticated-cyberattack-and-how-microsoft-defender-helps-protect/> [Accessed 13 February 2021].
13. Rosenbaum, R. (2012). Richard Clarke on Who Was Behind the Stuxnet Attack. Retrieved 18 March 2021, from https://www.smithsonianmag.com/history/richard-clarke-on-who-was-behind-the-stuxnet-attack-160630516/
14. US 'launched cyber-attack on Iran weapons systems'. (2019). Retrieved 18 March 2021, from https://www.bbc.com/news/world-us-canada-48735097

The Remote Workforce Will Lead to More Ransomware Incidents in 2021

continued from page 9

Zero Trust: A Zero Trust model, which includes safeguards such as multi-factor authentication and encrypted communications, is built on inherently not trusting any person, device, or network. Because of this, access to sensitive resources is earned. Zero Trust leverages identities and provides limited access to ensure that trusted identities get access to the applications, systems, networks, and data they are entitled to, based on their role and business or operational needs. Treating all communication and workloads as potentially hostile threats thwarts ransomware from moving laterally, limiting an incident from a single user to potentially an entire network. With Zero Trust in place, points of entry and unauthorized access are blocked because all network users are assumed to be threat actors. Institutions can therefore ensure business continuity without disruption through ongoing security, productivity, and compliance.

Prepare a response plan: The last thing an organization should do is wait to experience its first ransomware attack before making a plan to spot the indicators of compromise (IOCs). It's crucial to be prepared and have an incident response plan so the organization can rapidly fix the vulnerability. An automated incident response tool saves the security team time and is a better investment guaranteeing all prescribed steps are taken, and in the same order, ensuring nothing is missed. Furthermore, an automated incident response helps security teams bring relevant parties together as soon as a potential breach occurs to mitigate risk and reputational damage, which is crucial to maintain healthy business relations and trust.

Threat actors are still at large and seek to continuously implement new ransomware attack tactics in order to gain control of vital data and halt organizations' ability to operate. Hackers and organized cybercrime will continue to capitalize on any possible opportunity to disrupt an enterprise network. In addition to the above security controls, organizations should also implement and regularly update the cybersecurity awareness training for their employees to help decrease their chances of falling victim to easily avoidable threats, such as phishing. The training should teach them what to look for and how to respond should they suspect their device has been infected.

There were some hard lessons learned in 2020 as it relates to cybersecurity, with perhaps the biggest takeaway being that it is not just for large companies and cybersecurity should be appropriately funded across the board. Organizations need to have a solid strategy in place on how to face the proliferation of signals and separate those signals from noise as they manage risk. Adopting a Zero Trust architecture and increasing the centralization of visibility will be critical as the organization's ecosystem evolves through digital transformation and the growth of remote work.

Celebrating ISSA Women

continued from page 6

Still Work to Do (continued)

Statistics show 12% of ISSA working professionals rank in top executive level careers, with 39% representing senior level career status, and 17% representing mid-level careers. With 68% of ISSA membership ranking in mid to top level leadership positions, it would be interesting to see how this percentage breaks down in terms of women vs. men. I will tackle that in another column.

Conclusion

Through ISSA, we'll continue the momentum for women for leadership, member involvement and professional growth. As we look back at a year where resiliency and determination were key, ISSA women made inroads in leading the organization and the industry. Hats off to our ISSA women stepping up to make a difference.

About the Author

Dr. Curtis C Campbell, C|CISO, is VP of Atlantic Capital Bank in Atlanta, GA, serves as Director on the ISSA International Board, and is President of the ISSA Chattanooga Chapter. Curtis holds a Ph.D. in Organizational Leadership in Information Systems Technology, and serves on the advisory board of University of TN-Chattanooga, a national Center for Academic Excellence for Cyber-Defense (CAE-CD) studies. She was named ISSA Fellow

"Enjoying the Journal? Got ideas for how to improve it? Let us know by taking the official ISSA Journal Survey:

https://bit.ly/3powHrB

We want the Journal to reflect what you want and be a valuable part of your ISSA membership and security career."

Multi-cloud Security

By Pradeep Nambiar, CISSP Director of Technology, Chief Security Architect, Altran, Part of Cap Gemini

Use of multi-cloud strategies is increasing in business. Cybersecurity professionals need to be flexible to adapt to their use.

As cloud adoption in enterprises around the world is growing, we are starting to see a trend where enterprises are using more than one public cloud provider. A recent state-of-cloud report shows 93% of organizations are considering a multi-cloud strategy. Multi-cloud strategies helps to minimize downtime and disruptions for critical enterprise applications in the event of cloud provider outages. They also shield enterprises from cloud provider lock-ins and lead to better competitive pricing for procuring cloud services. Beyond the basic compute, network, and storage services, cloud providers are offering higher level cloud services such as Identity, Access & Single Sign-on management services, cloud security services, NOSQL database services, Analytics, AI and machine learning services, blockchain services, container services, storage services, cryptographic key management services, and more. These services are quickly becoming basic building blocks for enterprise cloud applications. Data governance, compliance, and user privacy regulations sometimes require data that is hosted by cloud with geographical region, country and state regulation considerations. All cloud provider services are not created equal and hence having a multi-cloud strategy helps enterprises optimize and choose the cloud services that best meets their need to gain a competitive edge in the marketplace.

Enterprise software vendors are engineering their applications so that they can be deployed in cloud environments, in addition to an on-premises. When offering cloud enabled services, vendors are realizing that these services must be able to deploy, and run, in multi-cloud environments such as Amazon Web Service, Azure, Google, IBM Cloud etc. Multi-cloud applications can be built to run completely within a single cloud provider but deploy to any public cloud provider, or a single cloud application can leverage cloud services and platforms from more than one cloud provider. When building to run within a single cloud provider, application developers would want to abstract common cloud services, including security, so that it can seamlessly deploy to any cloud with minimum configuration changes. Architecting multi-cloud applications to leverage the best of breed services from more than a single cloud provider must still build abstraction for cloud services used so that it can adapt to a new cloud provider in future or use intermediary vendors that support their APIs for multi-cloud.

Enabling applications for multi-cloud is becoming a differentiator and winning strategy for software vendors. Engineering software applications for multi-cloud offerings require selecting the right application platform, security considerations, and abstracting out common cloud environment differences and choosing cloud services carefully so that applications can be adapted to run seamlessly in any public cloud environment.

A multi-cloud security strategy that gives due consideration to the security at the time of engineering the application makes security much easier during the operational phase. This article describes how to address multi-cloud security for multi-cloud enabled applications and for enterprises leveraging services from multiple public cloud providers. Note: Security considerations are still valid in a hybrid cloud scenario also where enterprise may host certain application components on premise in a private cloud.

The heart of any cybersecurity program is to ensure availability, security, and protection of data. Moving applications to cloud increases the attack surface to the world wide web. Securing cloud-based applications also requires implementing a cybersecurity framework like an on-premise application and are generally modeled around a standard NIST cyber security framework that encompasses an Identify, Protect, Detect, Respond, and Recover strategy. The responsibility of protecting

the various cloud application layers becomes a shared responsibility between the cloud provider and the enterprise.

Image 1. – NIST Cyber Security Framework: https://www.nist.gov/cyberframework/framework

Cloud resources use virtual infrastructure that adds additional layers requiring protection. Cloud service can be subscribed using three basic service models, namely; Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). Depending on the nature of cloud service used to deploy the enterprise application, the security responsibilities between the enterprise and the cloud provider vendor varies. The following figure shows the shared responsibility model when it comes to cloud security between the cloud provider and the enterprise subscribing to the cloud service.

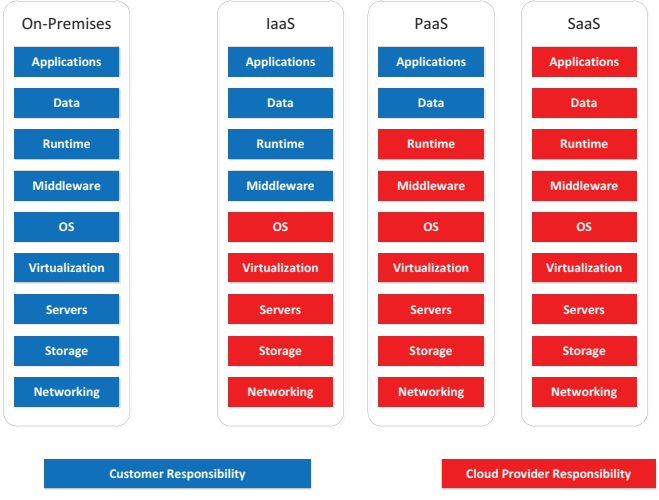

Image 2. Shared responsibility for security in the cloud

With the SaaS model, all responsibilities with regards to security is provided by the cloud provider, though due diligence from the consuming enterprise is a must to ensure that the cloud provider follows industry best practice and certifications for protecting enterprise data and that the SaaS provider supports cloud provider vendors that aligns with enterprises' multi-cloud strategy. Tenant isolation in SaaS is preferred from the perspective of data isolation and security and simplifies the SaaS provider's implementation of security policies to meet custom enterprise security policy requirements.

For IaaS and PaaS, the enterprise is responsible implementing a cyber security framework to protect areas which falls under their responsibility. To implement an effective cyber security framework the enterprise will have to understand the specific security services offered by the cloud provider. A multi-cloud strategy for consuming IaaS and PaaS must address operational security and security engineering during the cloud application development lifecycle as implementing security to protect the application and data will vary between cloud providers. DevSecOps, a security focused CI/CD/CM (Continuous Integration, Continuous Delivery and Continuous Management) pipeline tool for building, deployment and managing the cloud application security, must address and adapt to multi-cloud.

Challenges of Multi-cloud Security

A multi-cloud strategy must address complex security landscapes across various cloud providers with their own proprietary security implementations as it relates to authentication and access management. A shortage of skilled resources that are expert in a multi-cloud security makes it more challenging.

Securing access to resources in a multi-cloud environment requires careful consideration on a robust and scalable identity and access to management solutions that must support second factor authentication and are accessible from a host of different types of end points and a mobile workforce. Managing the keys, certificates, secrets, encryption used to protect data must be addressed as each cloud provider offers their own services for keys and certificate management.

Visibility and management of cloud resources across multiple cloud providers for monitoring and security compliance is another challenge that requires consideration.

To address multi-cloud security challenges we need to look at from the following perspectives –

- Multi-cloud Security Training Program
- Cloud Application Engineering Considerations
- Cloud Application Operational Considerations

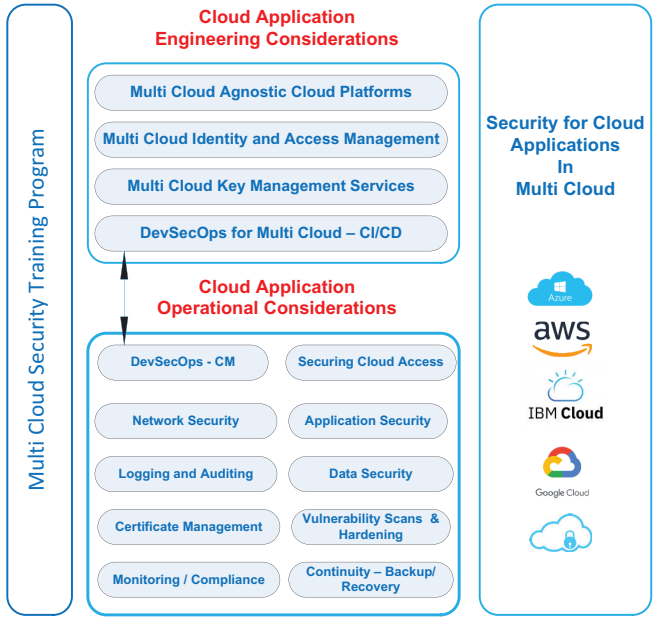

Figure 3 – Multi-cloud challenges

Multi-cloud Security

Multi-cloud Security Training Program

Applying security controls to a single cloud is hard. Applying security polices consistently across different public clouds requires skilled cloud security and Architects and Engineers that understand the nuances of the different public cloud providers. Enterprises will need to invest in training of technical resources for the various public cloud providers to manage the security controls effectively for multi-cloud.

All major cloud providers provide a comprehensive cloud security training program. Enterprises must set up a multi-cloud security training program to train their staff for both the engineering and operational services that the various cloud providers offer. Here is a list of cloud security training information from the major public cloud providers. Enterprises can complement these training with their own custom training based on their area of interest.

- Amazon Cloud – https://aws.amazon.com/training/learn-about/security/
- Microsoft Azure Cloud – https://docs.microsoft.com/en-us/learn/roles/security-engineer
- IBM Cloud - https://www.securitylearning-academy.com/
- Google Cloud - https://cloud.google.com/training

Cloud Application Engineering Considerations

Businesses enabling cloud applications for multi-cloud deployment need to consider the choice of security services the application will consume such as identity and access management, single sign-on service protocols, key management for data security, security configurations for application platform stack, and a DevSecOps enabled CI/CD pipeline that integrates security in to the deployment for better management in the operations of the application in multi-cloud environment.

Cloud Agnostic Application Platform

A natural choice for future proofing cloud applications to be multi-cloud is to develop or migrate applications using cloud native technologies. Cloud native application are a container-based environment that abstract the underlying compute, storage, and networking primitives without regard to the cloud provider. Container-based applications use an orchestrator to control and schedule application life cycle. The DevOps team manages resources and security policies defined via policy files. Kubernetes environment is available from all popular cloud vendors such as Amazon Elastic Kubernetes Service, Google Kubernetes Engine, Microsoft Azure Kubernetes Service, IBM Kubernetes Service etc. Securing and managing application workloads in a kubernetes environment requires additional components and work. IBM RedHat OpenShift is another platform that builds on the kubernetes foundation and adds an out of the box CI/CD pipeline, security and management stack for cloud native applications. Further, the availability of OpenShift on major cloud providers eases the deployment and management of cloud applications on any cloud.

Identity and Access Management

Identity and Access Management (IAM) is critical piece to any enterprise for managing authentication and authorizations to resources. Extending IAM to multi-cloud should be a key consideration in any multi-cloud strategy. Regardless of whether the enterprise opts to host the Identity provider on the premises or in the cloud, it must address authenticating users and support an access control policy that will work across multiple cloud providers. Most cloud provider support identity authentication and authorization and single sign on (SSO) support using protocols such as SAML (Security Assertion Markup Language), OAuth and OpenID Connect that allows delegating authentication to externally hosted identity and authentication and authorization servers. IAM solutions supporting multi-cloud must be robust, scalable and must support second factor authentications (2FA) for desktops, mobile, application APIs etc. A cloud based IAM can support and keep pace with the latest standards and SSO support while having the ability to federate with an on premise IAM. Consider adopting an IAM solution that can integrate with the all the cloud providers the enterprise chooses to incorporate in their multi-cloud strategy. Consider adopting a scalable IDaaS (Identity as a Service) hosted in the cloud that will support multi-cloud. Some examples of IDaaS are Azure AD, IBM Cloud Identity and Access Management, Okta etc.

Key Management Service

Encrypting data in the cloud requires encryption keys. Each cloud vendor offers their own key management service (KMS) with their own APIs which are vastly different. The cloud providers have now started supporting BYOK (Bring Your Own Keys) where keys can be owned by the enterprise outside the cloud infrastructure on-premises or at another location. Key management in a multi-cloud environment becomes more challenging and a well thought out strategy is required to manage keys to keep the key management vendor neutral and at the same time have complete control of the encryption keys. OASIS standard – KMIP (Key Management Interoperability Protocol) can address the challenges of interoperability. A KMIP compliant key management server can talk to a cloud provider Key Manager Service that are KMIP compliant. A KMIP compliant server stores, controls, and exchanges encryption keys, certificates, and secrets with clients (application) and servers. It is to be noted that popular cloud service providers like AWS, Azure and Google KMS do not support KMIP. Third party key management servers that are compliant with KMIP also provide plugins to interact with cloud provider KMS via KMIP. When choosing cloud virtualization hypervisors, applications and databases that require encryption consider KMIP support. Most popular databases such as DB2, MongoDB, MySQL etc. support KMIP for managing encryption keys and encrypting data. Developing applications that deal with encryption can use KMIP to keep they application portable across multi-cloud using a KMIP compliant key manager server solution.

DevSecOps for Multi-cloud – Continuous Integration and Continuous Delivery (CI/CD)

DevOps is an approach and technique to build, deploy and operate applications to cloud. Cloud Application Engineering considerations in DevSecOps deal with continuous integration (CI) and continuous delivery (CD). CI/CD must address the incorporation of security tasks such as static code scans, library vulnerabilities scan, dynamic testing scans etc. An efficient DevSecOps pipeline will continuously integrate, deploy, and manage operations of cloud applications. For multi-cloud support a DevSecOps pipeline can include multiple public cloud targets. Building a multi-cloud enabled DevSecOps pipeline can leverage the following libraries that abstracts services from multiple public clouds:

- Python Libraries for multi-cloud resources – https://libcloud.apache.org/
- Java Multi-cloud Toolkit - https://jclouds.apache.org/
- JavaScript Multi-cloud Library for Node.js - https://github.com/pkgcloud/pkgcloud
- Terraform - Creating Infrastructure as code to provision and manage multi-cloud - https://www.terraform.io/
- Cloud Foundry APIs for continuous delivery is an open source software with multi-cloud support. External services are accessed using Open Service Broker API - https://www.cloudfoundry.org

DevSecOps bridges the engineering and operational considerations. Cloud application operational considerations are described in the next section.

Cloud Application Operational Considerations

Security and compliance in a multi-cloud environment necessitates a consistent definition of security controls and policies that can be applied seamlessly across multi-cloud providers. Operations must also have continuous monitoring and visibility for managing security and compliance.

A multi-cloud strategy for an enterprise must at the minimum include implementation and usage guidance for the top cloud security controls as it relates to the various cloud provider an enterprise wants to engage with.

1. DevSecOps – Continuous Management (CM) – must address operational security requirements like enabling logging, dynamic security scans, continuous vulnerability scans and operational tasks for security governance.
2. Securing Cloud Access - Identity and Access Management (IAM) including Multi-Factor Authentication
3. Network Security - VPN, Firewall & ACLs management
4. Application Security - DDOS (Distributed Denial of Service) prevention and Web Application Firewall (WAF)
5. Logging and Auditing
6. Data Security – Data protection, Encryption and Key and secrets management
7. Certificates - Certificate Management
8. Vulnerability Scans / Hardening – Vulnerability scans and Patch/Update Management
9. Monitoring and Compliance
10. Operational Continuity - Secure Backup and Recovery

Enterprises must shortlist the cloud providers they want to do business with. Cloud security architects must evaluate the shortlisted cloud providers or third-party vendors from the cloud provider marketplace that offers services for each of the cloud security controls enumerated above. The implementation guidance must be documented for the engineering teams that develop and deploy multi-cloud applications. When evaluating, due consideration must be given to integration points within the cloud application. The shared responsibility model varies based on the nature of services the enterprise procures i.e., whether enterprise is deploying their own applications in a multi-cloud environment or whether the enterprise is procuring SaaS application services. Some of the services may be available via third parties which are generally available from the cloud provider marketplace.

The following table is a partial list of solutions and services available from the popular public cloud provider that addresses the top cloud security controls as currently available. Note: Cloud providers are continuously adding new services and this list may not reflect any new services since the writing of this document.

Top Cloud Security Controls	Amazon Cloud	Microsoft Azure Cloud	IBM Cloud	Google Cloud
DevSecOps CM	https://aws.amazon.com/devops/	https://azure.microsoft.com/en-us/services/devops/	https://www.ibm.com/cloud/devops	https://cloud.google.com/devops
Securing Cloud Access	AWS IAM, AWS Cognito, AWS SSO, AWS Directory Service, AWS Resource Manager, AWS Organization	Azure Active Directory	IBM Cloud IAM	Google Cloud IAM
Network Security	Security Groups and NACLs, AWS Firewall	Azure Firewall	IBM CIS, Firewall, Cloud Security Groups	Google Cloud Armor, Network Telemetry
Application Security, DDOS	AWS WAF, AWS Shield	Azure WAF	IBM CIS (Cloud Internet Service) - WAF	Google Cloud Armor
Logging / Auditing	AWS Cloudwatch	Azure Audit Logs	IBM Log Analysis with LogDNA	Google Cloud Logging
Data Security	AWS KMS, AWS Cloud HSM, Amazon Macie, AWS Secrets Manager, Third Party	Azure KMS, Azure Information Protect	IBM Key Protect, IBM Security Gaurdium	Google Cloud Key Management, Cloud DLP, Secrets Manager
Certificates	AWS Certificate Manager	App Service Certificate / Azure Key Vault	IBM Certificate Manager	Certificate Authority Service
Vulnerability Scans / Hardening	Amazon Inspector, AWS Trusted Advisor, AWS Systems Manager Patch Management	Azure Security Center, Azure Automation Update Management	Vulnerability Advisor, IBM Security Guardium VA, Third Party	Cloud Security Command Center, OS Patch Management
Monitoring / Compliance	AWS Security Hub, AWS Config, AWS Cloud Trail, AWS GuardDuty, Amazon Inspector	Azure Security Center - Azure Security Benchmark, Azure Policy, Third Party	IBM Cloud Security and Compliance Center, IBM Cloud Monitoring with Sysdig (for cloud native)	Cloud Security Command Center
Continuity - Backup/ Recovery	AWS Backup, Amazon S3 Glacier	Azure Backup, Azure Site Recovery	IBM Cloud Backup, IBM Cloud Object Storage	Cloud Storage - Nearline, Coldline, Archive

Conclusion

This article looked at some of the key drivers for enterprises that are accelerating a multi-cloud adoption, such as avoiding cloud provider lock in, leveraging best of breed cloud services, and competitive pricing. Multi-cloud adoption brings its own challenges to security. Breaking security from the perspective of cloud application engineering and cloud application management helps address the multi-cloud security challenges.

Multi-cloud vendor solutions and resource references

1. Cloud Report 2020 - https://info.flexera.com/SLO-CM-REPORT-State-of-the-Cloud-2020
2. https://jenkins-x.io/blog/2019/01/14/happy-first-birthday/#multicloud
3. https://www.ibm.com/cloud/garage/dte/tutorial/ibm-cloud-pak-multicloud-management-policy-management-hands-lab
4. https://www.vmware.com/cloud-solutions/multi-cloud-ops/security.html
5. Log DNA Multi-cloud https://logdna.com/blog/what-is-multi-cloud/
6. Multi-cloud security and compliance management -https://www.meshcloud.io/why-meshcloud/
7. Mapping on premise security controls and services available for major public cloud vendors - https://firegenanalytics.com/2019/02/mapping-of-security-controls-terminology-between-on-prem-and-public-cloud/
8. Thinking about IAM for multi-cloud: https://www.eweek.com/security/moving-to-multi-cloud-time-to-rethink-identity-access-management
9. Mesh cloud multi-cloud IAM : https://www.meshcloud.io/solution/cloud-identity-and-access-management/
10. OpenShift for multi-cloud: https://www.ibm.com/blogs/systems/how-red-hat-openshift-can-support-your-hybrid-multicloud-environment/
11. Key Management in multi-cloud: https://blog.equinix.com/blog/2019/03/05/key-management-in-multicloud-environments/
12. Cloud Foundry – Multi-cloud Service Brokers https://www.cloudfoundry.org/trainings/multi-cloud-cloud-foundry-making-real/
13. Key Management in Multi-cloud - https://www.helpnetsecurity.com/2020/04/20/byok/
14. Forrnetix solution for centralized key management for multi-cloud : https://www.fornetix.com/centralized-key-management-for-multi-cloud-architectures/
15. AWS and Microsoft Azure services equivalents: https://docs.microsoft.com/en-us/azure/architecture/aws-professional/services
16. Microsoft Hybrid and multi-cloud solutions: https://azure.microsoft.com/en-us/solutions/hybrid-cloud-app/
17. Multi-cloud DevOps Pipeline: https://techbeacon.com/enterprise-it/how-build-devops-pipeline-multi-cloud-app-deployment

INTEGRATE CYBERSECURITY TECHNOLOGY AND POLICY

Prepare to Confront Tomorrow's Threats with George Washington University's Online Master of Engineering in Cybersecurity Policy and Compliance Program

Shahram Sarkani, Ph.D., P.E.
Director of Engineering Management and Systems Engineering Online Programs

Adapting to cybersecurity threats has become an everyday challenge as attackers constantly identify new vulnerabilities in organizations ranging from state governments to financial services giants. Professionals can only get ahead of the next danger by gathering accurate intelligence and implementing data-driven solutions to guard essential systems and personal information.

Sophisticated technical solutions informed by predictive analytics are vital to address the evolving threats posed by data breaches, malware and hacks. But, as any cybersecurity professional knows, an organization's defenses rely just as much on workers and leadership as any security software. Well-informed policies and strict regulatory compliance are essential protections against compromising and potentially costly incidents.

An effective cybersecurity strategy covers all these bases with an agile, proactive approach. Technical professionals can lead the way by harnessing cutting-edge tools while establishing a culture of awareness and accountability.

George Washington University developed the online Master of Engineering in cybersecurity policy and compliance (M.Eng. [CPC]) to provide a multidisciplinary path toward a career in developing, implementing and enforcing best practices. This course of study—offered by an institution that the National Security Agency and the Department of Homeland Security recognized as a National Center of Academic Excellence in Cyber Defense Research—takes an engineering management

perspective on concepts in information security and computer science.

Students in the M.Eng.(CPC) program learn from GW's faculty of top researchers and engineers in courses exploring topics like:

- **Security systems**
- **Cyber incident forensics**
- **Cyberattack forecasting**
- **Situational awareness procedures**
- **Risk mitigation**
- **Cryptography**

Employers increasingly understand the importance of technical experts who can prepare their organizations to protect against data breaches and confront cyberattacks. That's why demand continues to grow for cybersecurity professionals: the Bureau of Labor Statistics projected a 32 percent increase in positions for information security analysts between 2018 and 2028, vastly outpacing the 5 percent average growth rate for all occupations. Joining GW's online graduate program in cybersecurity policy and compliance could help you advance your career while taking a stand against the next wave of threats.

Learn more at onlinecybersecurity.seas.gwu.edu/online-masters

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

Mitigating Attacks on a Supercomputer with KRSI

By Billy Wilson, billy_wilson@byu.edu

Supercomputer administrators face unique challenges securing their machines. This article looks at one tool to help overcome these challenges.

Abstract

Kernel Runtime Security Instrumentation (KRSI) provides a new form of mandatory access control, starting in the 5.7 Linux kernel. It allows systems administrators to write modular programs that inject errors into unwanted systems operations. This research deploys KRSI on eight compute nodes in a high-performance computing (HPC) environment to determine whether KRSI can successfully thwart attacks on a supercomputer without degrading performance. Five programs are written to demonstrate KRSI's ability to target unwanted behavior related to filesystem permissions, process execution, network events, and signals. System performance and KRSI functionality are measured using various benchmarks and an adversary emulation script. The adversary emulation activities are logged and mitigated with minimal performance loss, but very extreme loads from stress testing tools can overload a ring buffer and cause logs to drop.

Introduction

Systems administrators of high-performance computing (HPC) sites face the daunting task of securing research data without sacrificing peak system performance. They facilitate cutting-edge research that contractually comes with tight deadlines and stringent data security requirements. Satisfying both time and security constraints is an ongoing challenge that often requires novel approaches to old problems. This paper describes and tests Kernel Runtime Security Instrumentation (KRSI), a new Mandatory Access Control (MAC) extension in Linux. It allows systems administrators to program very specific and targeted MAC policies that potentially avoid the performance impact of large MAC extensions.

This author will refer to the technology as KRSI because it is a distinctive acronym that its creator continues to use in his presentations (Singh, 2020 July). However, the reader should be aware that this technology has been referred to as LSM BPF Hooks by Linux kernel developers (Corbet, 2019 December; Corbet, 2020) and LSM Probes in user-space applications (Olsa, 2020).

KRSI is a Linux Security Module (LSM) that hooks into the same kernel security events as SELinux and AppArmor, but rather than provide a major MAC extension, it lets an administrator compile and attach small, modular programs that control whether an action is allowed or denied (Singh, 2020 March). An administrator can attach their own custom code that controls file access, network activity, process execution, and much more.

This technology can potentially be adopted as an LSM of choice in high-performance computing. The fact that LSMs are disabled at HPC sites is prevalent enough that NIST included in their 2016 Action Plan Draft for HPC Security, "Consider why tools like SELinux don't get used" (National, 2016). Many systems administrators disable SELinux because of the negative performance impact it has on both synthetic benchmarks and real-world applications (Larabel, 2020).

Researching KRSI is a continuation of previous research on BPF Probes (Wilson, 2020 June). BPF Probes detected low-profile attacks against servers with little performance impact; however, the probes were limited in their ability to mitigate the attacks. In contrast, KRSI can provide both detection and mitigation.

In this research, new tracing scripts were written that used KRSI to detect and mitigate low-profile attack techniques. An environment of eight compute nodes was configured, booted from the latest available stable Linux kernel as of 21 September 2020. Benchmarking tools were run on the compute nodes to measure their baseline performance. A series of low-profile attacks were then launched, along with the tracing scripts, during a second set of benchmarks. Performance was compared and the functionality of the scripts was analyzed. Five appendices have been included that provide a KRSI tutorial, various source code, and benchmark results.

Technology Review

Writing KRSI programs is an advanced topic. To make the subject more approachable, a brief review of the technologies that KRSI is built upon is provided.

Berkeley Packet Filter (BPF)

KRSI is ultimately made possible by Berkeley Packet Filter, or BPF. Though traditionally recognized as a network filter tool, BPF is now a system-wide tracing subsystem for Linux.

Using BPF, systems administrators can write and attach small tracing programs to places of interest in the operating system. The programs can be attached to defined tracepoints or arbitrary functions, both in the kernel and user-space (Gregg, 2020). When a function is entered or exited, the BPF program can view the data passed to the function and data returned from the function. Though used primarily for performance analysis, BPF also serves as a valuable tool for security monitoring (Gregg, 2017).

A tracer called "bpftrace" was the tool of choice in previous research by Wilson (2020 June). It simplified the writing and attachment of BPF programs by providing an AWK-like syntax:

```
#!/usr/bin/bpftrace
probe1 /filter/ { action }
probe2, probe3 /filter/ { action }
```

Figure 1. Syntax of bpftrace

The three main components of bpftrace syntax are the probe, the filter, and the action. The probe specifies the tracepoint or function where the BPF program will be attached, the filter qualifies which events are processed, and the action defines the action to take when the event fires.

Security practitioners can leverage bpftrace to achieve a remarkable depth of visibility on Linux systems. Wilson (2020 June) provided bpftrace scripts to detect cryptocurrency software traffic, privilege escalation attempts, network pivot attempts, and SSH proxy creation.

Despite its broad monitoring capabilities, bpftrace was limited in its ability to mitigate attacks. At best, the tool could respond to an event by sending a signal to a process or by unsafely spawning a shell to perform an action (Gregg, 2020 September).

Wilson (2020 June) concluded that future research could focus on KRSI, an up-and-coming MAC extension that was better positioned to mitigate attacks with BPF.

Linux Security Modules (LSM)

Another essential prerequisite to writing KRSI programs is understanding how Linux Security Modules work. MAC extensions in Linux are implemented as LSMs, and this includes KRSI.

The LSM framework made it possible to extend the security model of Linux within the mainline kernel. Before its existence, Linux was limited to Discretionary Access Control, or DAC (Barkley, 1994). Projects that added MAC to Linux, such as Medusa, RSBAC, DTE, and the NSA's SELinux, had to maintain their own custom-patched kernels (Smalley, 2002 May).

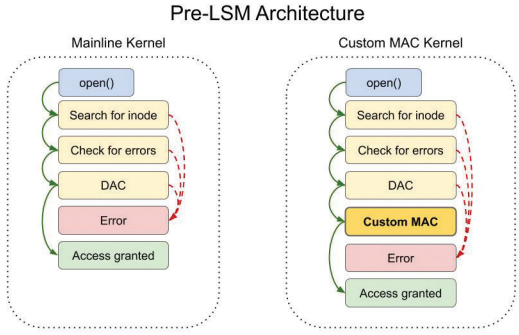

Figure 2. Pre-LSM Architecture that Required Custom Kernels for MAC

Linux kernel maintainers eventually created the LSM framework to provide a pathway for these custom security projects to be merged into the Linux mainline kernel (Smalley, n.d.). Multiple LSMs were eventually merged, but only one of them could be enabled at a time.

The two biggest players among Linux MAC extensions were SELinux for Red Hat-based distributions and AppArmor for Debian-based distributions (Smalley, 2002 June; Beattie, 2017). SELinux was known for type enforcement, which enforced how certain types of subjects could interact with certain types of objects. AppArmor took an alternate approach of basing its policies on filesystem paths.

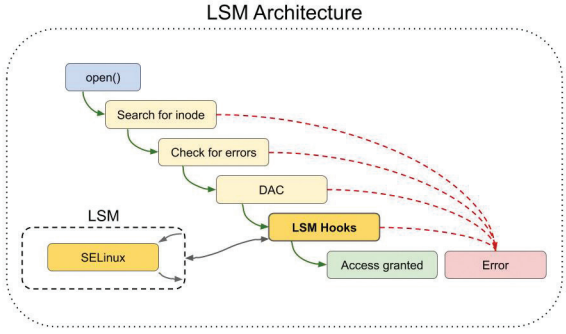

Figure 3. Linux Security Modules Architecture

There were several lesser-known MAC extensions as well, including Smack (Cook, 2017), TOMOYO (Takeda, 2009), and Yama (Cook, 2010). Because only one LSM could be enabled at a time, these smaller LSMs were often crowded out. However, starting in 2015, multiple LSMs could be loaded at the same time (Edge, 2015).

LSM hooks are still the common interface used by MAC extensions. They are listed in the Linux kernel source code file "/include/linux/lsm_hook_defs.h." The following are a few example entries:

```
LSM_HOOK(int, 0, inode_permission, struct inode
*inode, int mask)
LSM_HOOK(int, 0, bprm_check_security, struct
linux_binprm *bprm)
LSM_HOOK(int, 0, socket_listen, struct socket
*sock, int backlog)
```

Figure 4. Excerpt of /include/linux/lsm_hook_defs.h

The LSM_HOOK() macro specifies the function return type, the default return value, the name of the security hook, and the list of arguments passed into the hook.

The first line of the excerpt above is for a hook named inode_permission. Its arguments are an inode structure (which contains the metadata for a file) and an integer that represents the permission mask. When an LSM hook is triggered, control is passed to one or more LSMs which examine the arguments passed to the hook and then allow or deny access.

Another Linux kernel source file supplements information about these hooks. The following is an excerpt from "/include/linux/lsm_hooks.h" about the inode_permission hook, slightly modified for readability:

```
   * @inode_permission:
   *  Check permission before accessing an inode.
  ...
   *  @inode contains the inode structure to check.
   *  @mask contains the permission mask.
   *  Return 0 if permission is granted.
```

Figure 5. Excerpt of /include/linux/lsm_hooks.h

This entry documents that the inode_permission hook can be used for additional permission checks before allowing access to an inode. Security practitioners can reference these two source files to understand the purpose and usage of every LSM hook in the Linux kernel. They can be viewed online with the Elixir Cross Referencer at https://elixir.bootlin.com.

Kernel Runtime Security Instrumentation

In September 2019, KP Singh proposed a set of kernel patches to the Linux Kernel Mailing List for a new LSM called "Kernel Runtime Security Instrumentation," or KRSI (Singh, 2019 September). It allowed an administrator to attach BPF programs to the various LSM hooks, and it could also inject an error to block the operation in question. This gave administrators the ability to define their own MAC policies with arbitrary code.

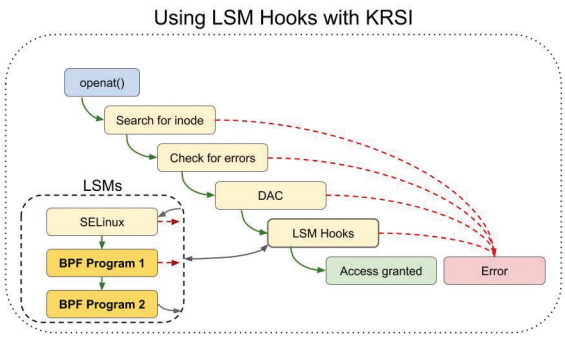

Figure 6. Using LSM Hooks with KRSI

The patches went through several revisions and were merged into the mainline Linux kernel in March 2020 (Borkmann, 2020). Two months later, Linux kernel version 5.7 was the first to include KRSI, released on 31 May 2020 (Corbet, 2020 June).

KRSI programs are being used in production at Google for mitigating various attacks, including LD_PRELOAD attacks (Singh, 2019 December).

Unfortunately, working examples of using KRSI in the mainline kernel are almost non-existent. The examples included in the KRSI patch set depended on special helper functions that were never merged into mainline. An open-source tool called Hawk demonstrated KRSI usage, but it only monitored process execution (Singh, 2020 September).

Tools for Writing KRSI Programs

There are three main toolsets for writing BPF programs: bpftrace, the BPF Compiler Collection (BCC), and the BPF-related tools provided in the Linux kernel source code.

The most approachable of the three tools is bpftrace, but it does not support KRSI yet. There is a pull request for this feature, but it has not been merged yet due to its dependency on a missing kernel helper function (Olsa, 2020). Once the pull request is merged, writing KRSI programs will become much simpler. The following is a basic example of using bpftrace to load a KRSI program that prevents any other BPF programs from being loaded:

```
bpftrace -e 'lsm:bpf { return -1234; }'
```

Figure 7. Example of using KRSI with bpftrace

This program attaches to the "bpf" LSM hook, which performs the initial check for all bpf() syscalls. It overrides the return value with a non-zero integer, causing all future attempts to call bpf() to fail until the program is unloaded. In other words, this BPF program blocks other BPF programs from loading.

The second tool of choice is BCC. Fortunately, it has supported KRSI since its 0.15.0 release in June 2020 (Song, 2020). As such, this research will rely on BCC.

Writing BCC scripts is significantly more involved than writing bpftrace scripts. There are two halves to each BCC script; the first half is the BPF program that will be loaded into kernel-space. This portion is written in C. The second half is the user-space script that will load the BPF program, poll data from it, and facilitate various command-line options. This portion is written in Python. The Python script can either embed the C program within it or reference it as a separate file. A full tutorial for writing a basic BCC script can be found in Appendix A.

Analysis of KRSI in HPC

The remainder of this paper is dedicated to measuring the ability of KRSI programs to mitigate attacks in an HPC environment. A test environment was built to compare the functionality and performance of KRSI-disabled and KRSI-enabled systems.

KRSI Scripts

Five BCC scripts were written that used KRSI for mandatory access control. Their source code is available in a public GitHub repo (Wilson, 2020 October). The scripts were written to address the following questions:

- Can KRSI block and report users who create files with insecure permissions?
- Can it block and report users who run unauthorized executables?
- Can it block and report users who establish SSH proxies?
- Can it block and report users who pivot to unauthorized network segments?

- Can it block and report unauthorized attempts to terminate processes?

All the scripts logged event data regarding the processes that caused them to fire. This data included the timestamp, command name, UID, GID, and PID of the processes, as well as any actions taken (allow or deny). Each script also recorded additional data that was unique to the type of event that it handled. The scripts wrote logs in a key-value format, but for the following sections, the logs were adjusted to a header-column format with a truncated timestamp for readability.

mac_fileperms

The mac_fileperms script was written to restrict users from setting the SUID (Set UID) and WOTH (Writable by Others) permission bits of files.

The SUID and WOTH bits are legitimate components of the Linux permissions model that are known to be abused by attackers. Setting the SUID bit on a file will allow a user to execute it with the file owner's privileges. It is necessary for executables like "passwd" and "sudo," but malicious programs have also set the SUID bit on files to maintain persistent backdoors (MITRE, 2020 August). The WOTH bit allows anyone to write to a given file. It is often set incorrectly by users to ensure an application works (MITRE, 2020 March) or to intentionally share data with peers. Malicious actors can find and abuse files that are writable by anyone.

The mac_fileperms script attached programs to the "inode_create" and "path_chmod" LSM hooks. These hooks were triggered when a new file was created or an existing file's permissions were modified, respectively.

The attached programs prevented SUID and WOTH permission bits from being set on new files or added to existing files. It did so by examining the requested permissions mode of a file and the umask of the process. If it detected a SUID or WOTH bit request, then it could prevent the file from being created or the file permissions from being updated.

Figure 8 below shows an invocation of the mac_fileperms script and its resultant effect on the user "billy," who has UID 1000.

Invocation of mac_fileperms:

```
# ./mac_fileperms -D -u billy
TIMESTAMP TYPE  COMM  UID  GID  PID   OLDMOD
REQMOD UMASK  NEWMOD ACT
T13:09:34 chmod chmod 1000 1000 18064 100664 -
-      004664 deny
T13:09:59 creat touch 1000 1000 18073 000000
100666 000000 100666 deny
T13:11:10 chmod chmod 1002 1002 18163 100664 -
-      004664 allow
T13:11:42 creat touch 1002 1002 18168 000000
100666 000000 100666 allow
```

User terminal:

```
billy@linux1 ~ $ touch /tmp/suidfile
billy@linux1 ~ $ chmod u+s /tmp/suidfile
chmod: changing permissions of '/tmp/suidfile':
Operation not permitted
billy@linux1 ~ $ umask
0002
billy@linux1 ~ $ umask 0000
billy@linux1 ~ $ touch /tmp/writable-by-others
touch: setting times of '/tmp/writable-by-
others': No such file or directory
billy@linux1 ~ $
```

Figure 8. Invocation and Effect of "mac_fileperms"

This script was invoked in Deny Mode, specifying that the user "billy" should be blocked from adding SUID or WOTH permission bits, whether through file creation or file modification. Other users were not restricted.

The user "billy" attempted to add SUID permission bits to a file, but chmod returned an error. The user then changed their shell's umask to include the WOTH bit for newly created files. When the user ran the touch command to create a file with the WOTH bit set, the command failed to create the file.

The script logged data from each event. The "chmod" type indicated a request to add SUID or WOTH permission bits to an existing file, and the "create" type indicated a request to set them on a new file. Both actions by the user "billy" were logged as denied.

These steps were then repeated with a different user. Their operations were examined and allowed, as evidenced by the two log entries for UID 1002.

mac_suidexec

The mac_suidexec script restricted SUID files from being executed. This contrasted with the previous script, which prevented their creation. Dozens of legitimate SUID executables exist on Linux systems, but centrally restricting SUID execution to a subset of files or users is desirable.

The script attached a program to the "bprm_check_security" LSM hook. This hook fired when a nascent process was being prepared for execution via the exec() family of syscalls. The hook exposed the linux_binprm structure, which contained information about the program being invoked (Drysdale). The invoked program's mode was examined for the SUID bit which, if detected, would result in the script logging the attempt and would possibly prevent the execution from occurring.

The following is an example of invoking mac_suidexec and its effect on a user session:

```
Invocation of mac_suidexec:
# ./mac_suidexec -D -u billy -F /bin/passwd
TIMESTAMP TYPE   COMM UID  GID  PID   DEV INODE
MODE   ACT
T14:48:15 exsuid bash 1000 1000 21307 42  43091310
104755 allow
T14:48:22 exsuid bash 1000 1000 21317 42  43091340
104111 deny
T14:48:30 exsuid bash 1000 1000 21325 42  43091580
104755 deny
User terminal:
billy@linux1 ~ $ passwd
Changing password for user billy.
Current password: ^C
billy@linux1 ~ $ sudo -i
-bash: /bin/sudo: Operation not permitted
billy@linux1 ~ $ newgrp
-bash: /bin/newgrp: Operation not permitted
```

Figure 9. Invocation and Effect of "mac_suidexec"

The above invocation of mac_suidexec prevented the user "billy" from invoking any SUID binaries on the system except for /bin/passwd (signified by the capital -F option). When the user attempted to run the "passwd" binary, it succeeded. However, attempts to run "sudo" and "newgrp" were blocked and logged by the script. The script logged the event type as "exsuid" and recorded the file's inode number, device number, and mode. The script tracked files by inode and device numbers rather than paths.

mac_sshlisteners

The mac_sshlisteners script prevented the creation of SSH proxy tunnels and handled both IPv4 and IPv6 addresses.

In "Securing the Soft Underbelly of a Supercomputer with BPF Probes," Wilson (2020 June) briefly explained TCP port forwarding. It is a built-in SSH feature that allows someone to use a server as a proxy to reach external resources. This feature can be used to transfer data to and from a device that lacks direct access to the internet. Although the feature can be disabled on SSH servers, it is on by default and usually left that way (Wilson, 2020 June). Whenever an SSH client attempts to listen on a socket, it indicates that the client is attempting to proxy traffic through a server that it can reach.

Unlike the BPF tracing script used by Wilson, which only detected SSH proxy tunnels, the mac_sshlisteners script both detected and prevented the proxy tunnels. It attached a program to the "socket_listen" LSM hook, which fired whenever a process attempted to change a socket to a LISTEN state. It examined IPv4 and IPv6 sockets to ensure that an SSH client was not attempting to listen on it.

The following was an invocation of the mac_sshlisteners script and its effect on a user who attempted to open an SSH proxy tunnel. Before invoking the script, the author slightly modified it to return the -NOLINK error upon denial instead of the -EPERM error. This was done to demonstrate that these scripts can return arbitrary error values to the user.

```
Invocation of mac_sshlisteners:
# ./mac_sshlisteners -D -U root
TIMESTAMP TYPE   COMM UID  GID  PID  PROTO
LADDR     LPORT ACTION
T08:53:39 listen ssh  1000 1000 6602 6      [::1]
9999  deny
T08:53:39 listen ssh  1000 1000 6602 6
127.0.0.1 9999  deny
User terminal:
billy@linux1 ~ $ ssh -D 9999 10.7.7.6
listen: Link has been severed
listen [::1]:9999: Link has been severed
listen: Link has been severed
listen [127.0.0.1]:9999: Link has been severed
channel_setup_fwd_listener_tcpip: cannot listen
to port: 9999
Could not request local forwarding.
Last login: Thu Oct  8 16:49:11 2020 from
192.168.100.15
billy@login1
```

Figure 10. Invocation and Effect of "mac_sshlisteners"

The script was invoked in Deny Mode with a capital "-U" to deny all users except root from opening proxy tunnels with SSH clients. The user then attempted to open an SSH proxy tunnel with dynamic port forwarding (signified by the SSH client's -D option). When the client attempted to change an IPv6 socket to a listening state on local port 9999, it was given the error that a link was severed. It then attempted another listen operation with an IPv4 socket, which failed in the same manner. The SSH connection to the remote host still succeeded for the user, but the proxy tunnel was not successfully established. The script logged both the IPv4 and the IPv6 attempts to change a socket to a listening state.

mac_skconnections

The mac_skconnections script restricted socket connections to certain destinations. It had full visibility into any socket connection attempt, but it was written to only examine IPv4 socket connections for this research. This script protected against pivot attempts by adding per-user firewall restrictions, and it did so without modifying the host's central firewall configuration.

The script was very similar to a BPF tracing script used by Wilson (2020 June) called tcp_connectfilter.sh. This script only handled TCP connections via the tcp_connect() kernel function, and it was limited to detection. In contrast, the mac_skconnections script handled any IPv4 socket connections regardless of the underlying protocol, and it provided both detection and mitigation.

The mac_skconnections script attached a program to the "socket_connect" LSM hook, which fired when a process attempted to make a socket connection. The following was an example of invoking the mac_skconnections script and how it affected a user.

Invocation of mac_skconnections:

```
# ./mac_skconnections -D 10.100.0.0 -m 255.255.0.0
-u billy
TIMESTAMP TYPE   COMM UID  GID  PID   PROTO
DADDR           DPORT ACT
T10:45:00 skconn ssh  1000 1000 10109 6
10.100.3.4      22    deny
T10:45:11 skconn ssh  1000 1000 10160 6
10.101.3.4      22    allow
T10:50:02 skconn dnf  0    0    10275 17
10.102.5.6      53    allow
T10:50:02 skconn dnf  0    0    10275 6
209.132.183.108 443   allow
T10:51:57 skconn nc   1002 1002 10357 6
10.100.10.11    80    allow
T10:57:20 skconn nc   1000 1000 10644 17
10.102.5.6      53    deny
```

User terminal:

```
billy@linux1 ~ $ ssh 10.100.3.4
ssh: connect to host 10.100.3.4 port 22:
Operation not permitted
billy@linux1 ~ $ ssh 10.101.3.4
Last login: Fri Oct  9 16:46:37 2020 from
192.168.10.15
billy@login3 ~ $ ^C
billy@linux1 ~ $ nc -u 10.100.4.5 53
Ncat: Operation not permitted.
```

Figure 11. Invocation and Effect of "mac_skconnections"

The script was invoked to block the user "billy" from making IPv4 socket connections to the 10.100.0.0/16 subnet. When the user attempted to SSH to a destination in that subnet, the operation was not permitted. Trying another destination outside the restricted subnet was successful. The script then logged some socket connections by the server's package manager, Dandified YUM (dnf). These connections were allowed. The user tried a final netcat to 10.100.4.5 on port 53, which was blocked by the script.

mac_killtasks

The mac_killtasks script restricted process signals. It was used to protect all the BCC scripts from early termination, whether by unprivileged users or root. It handled only SIGKILL and SIGTERM signals, but it could be extended to handle any signal.

Sending signals is a fundamental part of Linux and other POSIX operating systems. Signals can inform processes that an event has occurred (Kerrisk). They can also pause or terminate processes. The kernel generates them, but other system processes can request that the kernel send signals on their behalf. Over thirty signals are available on Linux, but the two restricted by this script were SIGTERM (ask a process to terminate itself) and SIGKILL (kill the process immediately).

The following is an example of invoking mac_killtasks to protect an instance of mac_fileperms and its effect on the root user.

Invocation of mac_killtasks:

```
# ./mac_fileperms -A -u root &
[2] 18290
# mac_fileperms_pid=$!
# ./mac_killtasks -D -e -t $mac_fileperms_pid
TIMESTAMP TYPE   COMM UID GID PID   TARGETUID
TARGETPID SIGNO ACT
T14:36:13 sgkill bash 0   0   11342 0
18290     15    deny
T14:36:17 sgkill bash 0   0   11342 0
18290     9     deny
T14:36:32 sgkill bash 0   0   11342 0
18316     9     deny
T14:36:44 sgkill bash 0   0   11342 0
18399     15    allow
```

A root terminal:

```
# pgrep -fl mac_
18290 mac_fileperms
18316 mac_killtasks
# kill 18290
-bash: kill: (18290) - Operation not permitted
# kill -9 18290
-bash: kill: (18290) - Operation not permitted
# kill -9 18316
-bash: kill: (18316) - Operation not permitted
# sleep 1000 &
[1] 18399
# kill 18399
```

Figure 12. Invocation and Effect of "mac_killtasks"

The script hooked into the "task_kill" LSM hook, which fired when a signal was about to be sent to a process. It examined the attributes of the source process and the target process to determine whether the signal should be allowed or not.

First, mac_fileperms was invoked to allow only the root user to create files with SUID or WOTH bits. The process ID of that script was saved. Then mac_killtasks was invoked so that no processes could send kill signals to the mac_fileperms process or to the mac_killtasks process that protected it.

In another terminal, the root user looked up the PIDs of the two scripts and attempted to send a SIGTERM signal to mac_fileperms. This attempt failed. The root user then unsuccessfully attempted to send SIGKILL signals to both BCC scripts. These also failed. Finally, root spawned a sleep process and sent it a SIGTERM signal. This succeeded and all activities were logged.

The script provided two new options: --kernel (-k) and --eternal (-e). The --kernel option modified the attached program to also control signals originating from the kernel itself in addition to those requested by user-space processes. This author did not test the implications of blocking signals from kernel-space processes and therefore marked this option as "dangerous" in the script's help text. The --eternal option ensured that the mac_killtasks executable itself was unkillable by any process except for its parent process. If the parent process exited, nothing could kill the process.

Test Environment

Two Linux kernels were compiled that differed only in whether KRSI was enabled or not. Version 5.8.10 of the source code was used, [1] which was the most recent stable version available as of 21 September 2020.

The Linux kernels were configured as similarly as possible to the kernel that is included in Red Hat Enterprise Linux distributions. This was not done manually, as the 5.8.10 kernel configuration file contains nearly 7,000 lines of options. Rather, an RPM package for the 5.8.10 kernel version was downloaded from ELRepo.org, [2] and the "config" file was extracted from it. Running `make oldconfig` with this configuration file returned no output, confirming that all options for the 5.8.10 kernel were defined.

The non-KRSI and KRSI kernels needed to be easily differentiated. One of the copies of the kernel source code was configured to append the string "-non-krsi" to its version. The other was configured to append the string "-krsi." This allowed for quick identification of the kernel in use by running `uname -r`.

KRSI was enabled in the latter kernel with the following configuration changes:

```
CONFIG_BPF_LSM=y
CONFIG_LSM="yama,loadpin,safesetid,
integrity,selinux,smack,tomoyo,apparmor,bpf"
CONFIG_DEBUG_INFO_BTF=y
```

Figure 13. Linux Kernel 5.8.10 Configurations to Enable KRSI

The first line enabled KRSI instrumentation. The second line provided a list of LSMs to initialize, with KRSI represented by the word "bpf" at the end of the string. The entries in the list besides "bpf" were included in the default configuration. The third line ensured that the kernel was compiled with BPF Type Format (BTF) symbols. Many BPF tools now depend on including BTF symbols in the kernel; the symbols help the in-kernel BPF Verifier perform memory access safety checks on the program before it is loaded. The pahole binary, part of the "dwarves" package, also needed to be installed to build the Linux kernel with BTF symbols.

The servers were likewise configured to be as similar as possible to each other. Eight compute nodes from an HPC cluster were reserved for testing. Kernel selection was handled with PXE boot. Each compute node mounted the same read-only root filesystem from a central NFS server.

Low-Profile Attack Script

A bash script simulated the activity that the BCC scripts were written to detect and mitigate. Every one-to-fifteen seconds, the script randomly performed one of the following actions:

- Created a new other-writable file
- Added other-writable permissions to an existing file
- Added the SUID bit to an existing file
- Ran a SUID executable file
- Attempted to open an SSH proxy connection
- Attempted to connect to a remote destination over TCP
- Attempted to connect to a remote destination over UDP
- Attempted to terminate the KRSI scripts (as root)

This script logged all actions taken, which allowed for comparison between the logs of the attack script and the logs of the detection scripts. The script is included in Appendix B.

Benchmarks

High-performance computing applications vary immensely in how they exercise a system. They can cause performance bottlenecks in processing, memory operations, filesystem operations, network communications, and more.

For this reason, three benchmarks were chosen for this research: xhpl for computational benchmarking, a C program named mdstress for IO benchmarking, and tcpkali for network benchmarking. These benchmarks are described below.

Computational Benchmark: xhpl

The High-Performance Linpack Benchmark, or xhpl, measures the computational performance of the largest supercomputers in the world. It solves a series of linear algebra equations to measure the maximum "flops," or floating-point operations per second, that a cluster is capable of. It can measure the flops of a single compute node or multiple compute nodes working in unison.

The xhpl benchmark was run on the eight test nodes individually. For each node, it ran thirty times on the non-KRSI kernel, thirty times on the KRSI kernel without the BCC scripts loaded, and thirty times with all five BCC scripts loaded.

Filesystem Benchmark: mdstress

The mdstress benchmark is a rudimentary C program written by this author. It created a new file, wrote the string "mdstress" to it, and deleted the file in a tight loop. When all loops completed, it printed the total elapsed time in seconds. The purpose of this benchmark was to determine how mac_fileperms behaved under extreme load. The source code of the program can be found in Appendix C.

Each mdstress benchmark was configured to complete in 10 seconds at optimal performance. Any overhead would cause its completion to exceed 10 seconds. For example, when creating 100,000 inodes, the stresser loop was rate limited to 10,000 loops per second, which at its fastest would complete in 10 seconds. Any slowdown beyond the rate limit would result in the benchmark taking longer than 10 seconds to complete. Each configuration was run thirty times on each node with no BCC scripts running, and the results were averaged. These tests were repeated with mac_fileperms running and then repeated one more time with the umask of the mdstress process set to 0000, which resulted in the WOTH bit being set on the inodes. File and directory caches were dropped before each run to minimize the effect of caching between runs.

Network Benchmark: tcpkali

The tcpkali benchmark can establish and tear down thousands and even millions of TCP connections in a short period of time. It is used to stress-test applications and networks. For this research, it helped identify the performance cost of attaching programs to the socket_connect LSM hook.

This benchmark used one compute node as a client. It used up to seven other computes nodes for destinations. The compute nodes were all tuned to allow up to 55,000 TCP connections per peer. These tunings can be found in Appendix D.

The client attempted 1,500 TCP connections per second per destination, scaling up to 10,500 TCP connections per second when all seven destinations were in use. Each run was given a thirty-second time limit, and the number of successful connections was recorded by the nodes acting as destinations. The destination nodes ran `socat` to listen for and record TCP connections. These tests were run without mac_skconnections loaded and with it loaded to compare performance.

Results

The benchmark results showed that the BCC scripts performed very favorably under intensive but realistic workloads. However, when mdstress and tcpkali pushed systems to very extreme levels, enough to degrade general system performance, then running the BCC scripts worsened performance. Those extreme workloads also overloaded the "perf ring buffer," which was the buffer used by BPF to stream high volumes of kernel events to user-space. This did not impede the BCC scripts in preventing attacks, but it did affect their ability to log them.

Detection Results

During the xhpl benchmarks, the adversary emulation script and all the BCC scripts ran and logged their activities. These logs were aggregated to determine whether the actions by the adversarial script were successfully detected and mitigated.

All malicious file activities, including creating world-writable files and adding WOTH or SUID bits, were logged as "denied" except for one of the 2,964 "Add WOTH bit" actions. All attempts to run SUID executables were denied and logged. All attempts to open SSH proxies were logged as "denied" except for two of the 2,943 IPv4 attempts. All unauthorized socket connection attempts and all attempts to kill the BCC scripts were denied and logged.

It was difficult to determine the cause of the three missing logs due to a shortcoming of the attack script itself. The attack script logged what actions it would take, but it did not record whether the attempted action was executed or not. It was possible that a few malicious actions failed from regular system errors. This would have acted as a short-circuit prior to the LSM hook being triggered and the BCC script coming into play.

The results are summarized below in Figure 14.

Malicious Action	Count	KRSI Script	Deny Count	Deny Types	Type Count
Create WOTH	2997	mac_ fileperms	8907	create	2997
Add WOTH	2964			chmod (0602)	2965
Add SUID	2946			chmod (4600)	2946
Run SUID	2989	mac_ suidexec	2989	exsuid	2989
SSH Proxy	2945	mac_sshlis-teners	5888	IPv4	2943
				IPv6	2945
Socket Connect	2918	mac_ skconnect	2918	IPv4	2918
Kill BCC (fileperms)	586	killtasks (file)	1176	fileperms only	586
Kill BCC (suidexec)	602	killtasks (suid)	1192	suidexec only	602
Kill BCC (ssh...)	600	killtasks (ssh...)	1190	sshlisten-ers only	600
Kill BCC (skconnect)	621	killtasks (skc...)	1211	skcon-nect only	621
Kill BCC (killtasks)	590				

Figure 14. Table of Malicious Action Counts and MAC Denial Counts

While the xhpl detection results were very favorable, the tcpkali and mdstress benchmarks demonstrated that an excessive amount of MAC events could result in dropped logs. These details will be discussed as part of the performance results.

Performance Results

Three tools measured the performance impact of KRSI: xhpl, mdstress, and tcpkali. While xhpl was configured to use a large but reasonable portion of each system's computational resources, mdstress and tcpkali scaled IO and network loads past reasonable system capabilities. Stress testing IO and network resources in this manner revealed the behavior of KRSI on systems suffering from extreme loads.

This section presents the benchmark results as charts; the numeric tables used to generate these charts can be found in Appendix E.

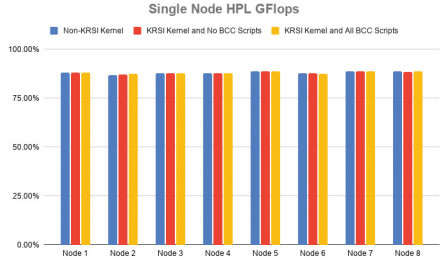

Figure 15. Chart of "xhpl" Results

KRSI had a less than 1% computational impact on xhpl benchmarks on the compute nodes. The worst case was Node 6, which exhibited a 0.14% performance loss when using a KRSI-enabled kernel with all BCC scripts loaded. Oddly, Node 2 showed a 0.51% gain in performance with BCC scripts loaded, but this would not be explained by KRSI. It was likely due to other factors in the test environment, such as network contention from unrelated compute nodes or system jitter. The rest of the nodes showed a performance change of 0.05% or less between the non-KRSI kernel and the KRSI kernel with BCC scripts running.

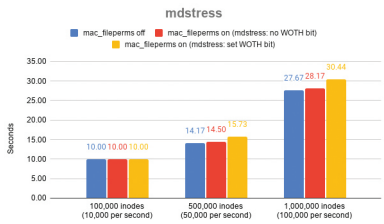

Figure 16. Chart of "mdstress" Results

The mdstress benchmark revealed two issues when under extreme filesystem load. First, if a compute node was stressed with more file creation operations than it could handle, then running BCC scripts that monitored file creation would make the problem worse. Second, when events were generated by the kernel-space program too quickly, the perf ring buffer would overwrite the oldest event data before user-space scripts logged it. This was indicated by messages from the BCC script stating, "Possibly lost N samples," with N ranging from 1 to over a million, depending on the load that mdstress placed on the script.

When handling 10,000 inode creation events per second, the mac_fileperms script did not cause performance loss issues, but it did drop logs due to the overwhelmed ring buffer.

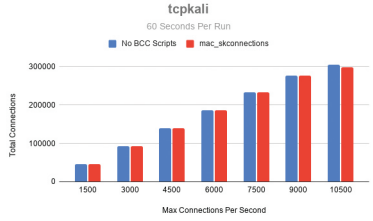

Figure 17. Chart of "tcpkali" Results

The tcpkali benchmark stressed the system beyond its ability to maintain its connection rate when set to 9,000 TCP connections per second. This performance degradation became worse at 10,500 TCP connections per second. The mac_skconnections script began to lose logs from the perf_ring_buffer at this point. In both cases, attaching mac_skconnections caused the system to perform slightly worse.

Conclusions and Future Work

Running BCC scripts with KRSI did not cause performance loss unless the system was already suffering from degraded performance. It did not impact xhpl benchmarks. It only impacted mdstress and tcpkali benchmarks after extreme loads were already causing general system performance problems. Therefore, companies and organizations should continue investigating KRSI for adoption in HPC and information security.

The BCC scripts appeared to still mitigate attacks even under extreme load, but they did not guarantee zero loss of logs. When the filesystem event rate approached the order of 100,000 file creations per second, the mac_fileperms script dropped millions of incoming logs. This is not a shortcoming of KRSI, but rather is the expected consequence of overloading a ring buffer. It is possible to increase the perf ring buffer's size, but the default is already 64 pages per CPU (Drayton). Any load that overwhelms the perf ring buffer has already pushed the system well beyond reasonable limits.

KRSI will become much more approachable once bpftrace supports it. This tool provides an AWK-like syntax that is much more intuitive to systems administrators and information security practitioners than the complex BCC scripts written for this research.

Those interested in KRSI can now experiment with it using the stock kernel of a major Linux distribution. Canonical released the 20.10 version of Ubuntu on 22 October 2020, and it runs on the 5.8 Linux kernel. Users can enable KRSI by specifying "bpf" in the "lsm" kernel parameter (e.g., "lsm=lockdown,yama,integrity,apparmor,bpf").

This author used a kernel configuration baseline that had many LSMs initialized, which may have drowned out the true performance impact of enabling KRSI. Future research can measure the performance of systems with either no LSMs initialized or only KRSI initialized.

The scripts written for this research used only six LSM hooks, available in the 5.8.10 Linux kernel. It would be valuable for those with kernel development and information security experience to write additional BCC scripts that leverage these LSM hooks in interesting new ways.

Once KRSI is generally available, HPC environments can potentially couple it with their schedulers to launch custom MAC policies per compute node, based on the users' jobs dispatched to those nodes. This would result in special security measures that follow users to whichever nodes their research is running on.

The ultimate benefit of KRSI is the freedom it gives systems administrators to build creative new MAC policies. Those who work on Linux systems with recent kernels are encouraged to try KRSI, be creative in its usage, and share their innovations with the community.

References

1. https://git.kernel.org/pub/scm/linux/kernel/git/stable/linux.git/tag/?h=v5.8.10
2. Mirror: https://linux.cc.iitk.ac.in/mirror/centos/elrepo/kernel/el7/x86_64/RPMS/
3. Barkley, J. (1994). Discretionary Access Control. NIST Special Publication 800-7. Retrieved 26 September 2020 from http://ftp.gnome.org/mirror/archive/ftp.sunet.se/pub/security/docs/nistpubs/800-7/main.html
4. Beattie, S. (2017). About AppArmor. AppArmor Security Project Wiki. Retrieved 26 September 2020 from https://gitlab.com/apparmor/apparmor/-/wikis/About
5. Borkmann, D. (2020). Merge branch 'bpf-lsm' [Computer Software]. Kernel.org git repositories. Retrieved 26 September 2020 from https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=641cd7b06c911c5935c34f24850ea18690649917

6. Cook, K. (2010). security: Yama LSM [Computer Software]. LWN. Retrieved 26 September 2020 from https://lwn.net/Articles/393012/
7. Cook, K. (2017). Smack. The Linux kernel user's and administrator's guide. Retrieved 26 September 2020 from https://www.kernel.org/doc/html/v4.15/admin-guide/LSM/Smack.html
8. Corbet, J. (2019, October). BPF at Facebook (and beyond). LWN. Retrieved 26 September 2020 from https://lwn.net/Articles/801871/
9. Corbet, J. (2019, December). KRSI – the other BPF security module. LWN. Retrieved 26 September 2020 from https://lwn.net/Articles/808048/
10. Corbet, J. (2020, June). The 5.7 kernel is out. LWN. Retrieved 26 September 2020 from https://lwn.net/Articles/821829/
11. Drayton, M. (2017, February). Make perf ring buffer size configurable [Computer software]. Github. Retrieved 21 October 2020 from https://github.com/iovisor/bcc/pull/997
12. Drysdale, D. (2015). How programs get run. LWN. Retrieved 17 October 2020 from https://lwn.net/Articles/630727/
13. Edge, J. (2015). Progress in security module stacking. LWN. Retrieved 26 September 2020 from https://lwn.net/Articles/635771/
14. Gregg, B., et al. (2016). Bcc Python Developer Tutorial. Github. Retrieved 1 October 2020 from https://github.com/iovisor/bcc/blob/master/docs/tutorial_bcc_python_developer.md
15. Gregg, B. (2018, September). gethostlatency.bt [Computer software]. Github. Retrieved 26 September 2020 from https://github.com/iovisor/bpftrace/blob/master/tools/gethostlatency.bt
16. Gregg, B. (2020, January). BPF Performance Tools: Linux System and Application Observability. United States: Addison-Wesley.
17. Gregg, B., et al. (2020, August). bcc Reference Guide. Github. Retrieved 26 September 2020 from https://github.com/iovisor/bcc/blob/master/docs/reference_guide.md
18. Gregg, B., et al. (2020, September). bpftrace Reference Guide. GitHub. Retrieved 26 September 2020 from https://github.com/iovisor/bpftrace/blob/master/docs/reference_guide.md
19. Gregg, B., & Maestretti, A. (2017, February). Security Monitoring with eBPF. In BSidesSF 2017, San Francisco, CA. Retrieved 26 September 2020 from https://www.youtube.com/watch?v=44nV6Mj11uw
20. Kerrisk, M. (2010). The Linux Programming Interface: A Linux and UNIX system programming handbook. San Francisco: No Starch Press.
21. Larabel, M. (2020). The Performance Cost to SELinux on Fedora 31. Phoronix. Retrieved 14 November 2020 from https://www.phoronix.com/vr.php?view=28798
22. National Institute of Standards and Technology. (2016). An Action Plan for High Performance Computing Security, Working Draft. Gaithersburg, MD. Retrieved 26 September 2020 from https://www.nist.gov/system/files/documents/2018/03/15/working_draft_actionplanhpc.pdf
23. Olsa, J. (2020). Add lsm probe support [Compute Software]. Retrieved 26 September 2020 from https://github.com/iovisor/bpftrace/pull/1347
24. Singh, K. (2019, September). Kernel Runtime Security Instrumentation. LWN. Retrieved 26 September 2020 from https://lwn.net/Articles/798918/
25. Singh, K. (2019, December). lsm_audit_env.c [Computer Software]. Github. Retrieved 1 October 2020 from https://github.com/sinkap/linux-krsi/blob/patch/v1/examples/samples/bpf/lsm_audit_env.c
26. Singh, K. (2020, March). MAC and Audit Policy using eBPF (KRSI). LWN. Retrieved 26 September 2020 from https://lwn.net/Articles/815826/
27. Singh, K. (2020, July). KRSI (BPF + LSM) – Updates and Progress. In Linux Security Summit North America, Virtual Conference. Retrieved 26 September 2020 from https://ossna2020.sched.com/event/ckpL/krsi-bpf-lsm-updates-and-progress-kp-singh-google
28. Singh, K., et al. (2020, September). hawk [Computer Software]. Github. Retrieved 17 October 2020 from https://github.com/googleinterns/hawk
29. Smalley, S., et al. (n.d.). Linux Security Modules: General Security Hooks for Linux. The Linux Kernel Archives. Retrieved 26 September 2020 from https://www.kernel.org/doc/html/latest/security/lsm.html
30. Smalley, S., et al. (2002, June). Linux Security Module Framework. In Ottawa Linux Symposium, Ottawa, Ontario, Canada. Retrieved 26 September 2020 from https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.119.1604&rep=rep1&type=pdf
31. Smalley, S., et al. (2002, May). Implementing SELinux as a Linux Security Module. Retrieved 26 September 2020 from https://www.nsa.gov/Portals/70/images/resources/everyone/digital-media-center/publications/research-papers/implementing-selinux-as-linux-security-module-report.pdf
32. Song, Y. (2020). prepare for release v0.15.0 [Computer Software]. GitHub. Retrieved 26 September 2020 from https://github.com/iovisor/bcc/commit/e41f7a3be5c8114ef6a0990e50c2fbabea0 e928e#diff-45c17c0a080fbe29e2b8ded8940aa1e8
33. Takeda, K. (2009). TOMOYO Linux Overview. Presented at linux conf au, Hobart, Australia. Retrieved 26 September 2020 from https://osdn.net/projects/tomoyo/docs/lca2009-takeda.pdf
34. MITRE. (2020, March). Keydnap. MITRE ATT&CK. Retrieved 17 October 2020 from https://attack.mitre.org/software/S0276/
35. MITRE. (2020, August) CWE-732: Incorrect Permission Assignment for Critical Resource. Common Weakness Enumeration. Retrieved 17 October 2020 from https://cwe.mitre.org/data/definitions/732.html
36. Wilson, B. (2020, June). Securing the Soft Underbelly of a Supercomputer with BPF Probes. SANS Institute. Retrieved 26 September 2020 from https://www.sans.org/reading-room/whitepapers/linux/securing-soft-underbelly-supercomputer-bpf-probes-39635
37. Wilson, B. (2020, October). bcc-lsm-scripts [Computer Software]. Github. Retrieved 17 October 2020 from https://github.com/wilsonwr/bcc-lsm-scripts

Upcoming Webinar

HOW TO BUILD IN FLEXIBILITY TO CREATE A MORE EFFECTIVE SECURITY STRATEGY

April 14th | 1PM ET

Ray Espinoza
CISO
Cobalt

Jack Roehrig
CISO
Turnitin

Register Now

ISSA
Information Systems Security Association

The Cyber Executive Forum is a peer-to-peer event – Members can feel free to share concerns, successes, and feedback in a peer-only environment.

ISSA Cyber Executive Membership Program

The role of information security executive continues to be defined and redefined as the integration of business and technology as it evolves. While these new positions gain more authority and responsibility, peers must form a collaborative environment to foster knowledge and influence that will shape the profession.

The Information Systems Security Association (ISSA) recognizes this need and created the exclusive Cyber Executive Membership program to give executives an environment to achieve mutual success. Connecting professionals to a large network of peers, valuable information, and top industry experts the program is a functional resource for members to advance personal and industry understanding of critical issues in information security.

Membership Benefits

- Free registration at four Cyber Executive Forums per year, including lodging for one night and all meals at each Forum
- You'll be part of an effective forum for understanding and influencing relevant standards and legislation
- Extensive networking opportunities with peers and experts on an ongoing basis
- Direct access to top subject matter experts through educational seminars
- CPE credits you earn will be automatically submitted
- Vendor Influence: A unified voice to influence industry vendors
- Online Community: Privileged access to our online community

Visit Cyber Executive Forum for more information or to register for the Forum.

Asia Pacific
Bangladesh
Chennai
Dehradun
India
Philippines

Canada
Alberta
Ottawa
Quebec City
Vancouver

Europe
Brussels European
France
Germany
Italy
Netherlands
Poland
Romania
Spain
Switzerland
Turkey
UK
Ukraine

Latin America
Argentina
Barbados
Bolivia
Brasil
British Virgin Islands
Chile
Colombia
Ecuador
Peru

Middle East
Bahrain
Egypt
Iran
Israel
Kazakhstan
Kuwait
Qatar
Saudi Arabia

USA
Alamo San Antonio
Blue Ridge
Boise
Buffalo Niagara
Capitol of Texas
Central Alabama
Central Florida
Central Indiana
Central Maryland
Central New York
Central Ohio
Central Plains
Central Texas
Central Virginia
Charleston
Charlotte Metro
Chattanooga
Chicago
Colorado Springs
Columbus
Connecticut
Dayton
Delaware Valley
Denver
Des Moines
East Tennessee
Eastern Idaho
Eugene
Fayetteville/Fort Bragg
Fort Worth
Grand Rapids
Grand Traverse
Greater Augusta
Greater Cincinnati
Greater Spokane
Hampton Roads
Hawaii
Inland Empire
Kansas City
Kentuckiana
Kern County
Lansing
Las Vegas
Los Angeles
Metro Atlanta
Mid-South Tennessee
Middle Tennessee
Milwaukee
Minnesota
Motor City
National Capital
New England
New Hampshire
New Jersey
New York Metro
North Alabama
North Dakota
North Oakland
North Texas
Northeast Florida
Northeast Indiana
Northeast Ohio
Northern Colorado
Northern
Virginia (NOVA)
Northwest Arkansas
Northwest Ohio
Oklahoma
Oklahoma City
Orange County
Phoenix
Pittsburgh
Portland
Puerto Rico
Puget Sound (Seattle)
Quantico
Rainier
Raleigh
Rochester, NY
Sacramento Valley
San Diego
San Francisco
Silicon Valley
South Bend - Michiana
South Florida
South Texas
Southeast Arizona
Tampa Bay
Tech Valley Of New York
Texas Coastal Bend
Texas Gulf Coast
Triad of NC
Upstate SC
Utah
Ventura County
West Texas
Wyoming
Yorktown

ISSA.org => Chapters