Automating Cyber / IT Processes for the Modern Work Force, A Post COVID-19 World

A WHITE PAPER

Date: 29 December 2020

By

MARKESMAN GROUP

11817 Canon Blvd, Suite 610
Newport News |VA 23606
TEL: (609) 558-2898
Point of Contact: Alex Wang
Alex.Wang@Markesman.com
CAGE Code: 75RH6 | DUNS 079459870

AUTOMATING CYBER RISK ANALYSIS

Recent advances in computer and network technology have led to substantial efficiency and performance gains for government entities across the globe. These advances have allowed remote offices to operate as increasingly integrated components of headquarters locations, supply chain and Defense Industrial Base (DIB) partners to better integrate directly with government operations and empowered remote personnel to remain productive even when on the road.

Unfortunately, the substantial productivity gains offered by these technological advances is often overshadowed by the time and labor-intensive process of maintaining cyber security standards in an ever-changing technological environment. Often the cost of a single, significant government computer system compromise can severely harm national security posture and greatly degrade public trust in government organizations.
Most government agencies and partners have retained cybersecurity experts to implement active cybersecurity measures and remediate new and discovered vulnerabilities. While this is a crucial step in securing cyber infrastructure, it is still a constant challenge for these teams to determine where to dedicate their limited time and resources within increasingly large and complex organizations. Additionally, these teams rarely have robust insight into the cyber risks presented to the government by third party commercial partners who have been trusted with sensitive government network access and information.

Automated cyber risk management is a similar yet distinct discipline from cybersecurity and aims to present an automated evaluation of the cybersecurity risk presented by different elements of a single organization and their third-party partners. This data-driven analysis allows the limited cybersecurity resources available to an organization to be deployed in a way that most efficiently limits the risk of a potential compromise. Within commercial organizations the risk associated with various vulnerabilities is usually presented in terms of potential financial cost to the target company. This methodology can be tailored to the needs of government organizations and their partners by incorporating the negative impact to the organization's mission effectiveness in addition to the potential financial toll.

The challenge of performing large scale cyber risk analysis can be overcome by automating the collection, analysis and potential impact scoring of internal and external security information and event management (SIEM) information, openly available internet scanning data, infrastructure and domain blacklists and vulnerability databases to present the end user with a comprehensive visualization of the automated cyber risk score currently assigned to each element of their organization. Automated analysis of the inherent risk associated with each event can be performed through machine learning, allowing the cybersecurity analyst to focus their valuable time on vulnerability remediation with priority already assigned to the highest risk areas of the organization.

GENERATING MEANINGFUL CYBER RISK SCORES

The most commonly used method for rating cyber risk is through the use of a numerical "risk score." Use of a numerical risk score, generally rendered as values between 0 and 100, allows cyber analysts to quickly assess the severity of identified risk factors and allocate resources to mitigating the risk as efficiently as possible.
These numerical scores are applied to discrete network segments that the customer would wish to monitor. For example, if a company had a headquarters and two satellite offices each might be assigned independent risk scores in order to track risk to each network segment independently. In other scenarios, such as when tracking the risk of cyber compromise to a third-party company within a larger portfolio of companies within a supply chain network, the company as a whole may be assigned a risk score.

It is also helpful to assign sub-scores to different risk factor categories for each overall risk score. This allows for more granular identification of the most serious types of cyber risks impacting the overall score. For example, the overall risk score for a company may be 73 out of 100 but the sub factors indicate that the negative impact to the overall score is largely due to the IT hygiene sub-factor. This would lead indicate to an analyst that resources would be best spent mitigating general issues related to IT hygiene rather than other potential subfactors such as email security or software vulnerabilities.

RISK FINDINGS

While the risk scores are the aggregated overview of numerous internal and external risk factors, each score and sub-score is itself made up of an analysis of various risk findings. Each risk finding is a discrete observation that indicates potential risk. When the aggregated together by the risk modeling algorithm all of the numerous findings are used to generate an overall risk score.
Each finding should correspond to a particular risk category, with its own associated sub-score, and these categories are generally weighted to indicate the overall severity of the type of risk. This weighting prevents the overall score from being skewed by high frequency but low severity findings. For instance, five email servers with improperly implemented DKIM protections would be weighted lower than a single instance of a mail server being vulnerable to a critical remote code execution vulnerability.

ITERATIVE ANALYSIS

In order to provide a complete understanding of the cyber risk landscape for a given network, scoring and analysis must be done continuously and iteratively. Providing a risk score for only a single point of time is significantly less valuable and actionable for cyber responders than providing continuous risk monitoring.

Analyzing changes to risk patterns over time allows a larger set of findings to be aggregated into the overall risk score. This in turn results in a more holistic view of the risk landscape presented by a given network. Additionally, it allows for intermittent or otherwise temporary risks to be better captured and presented to the risk analyst. Iterative analysis of risk factors over periods of time also provides greater scoring accuracy by mitigating temporary data collection issues. If access to a key threat intelligence data feed is lost for a window of time that temporary deficiency is easier to identify if monitoring is continuous over a longer timeframe. If the system is only developing a risk score for a specific point in time, then the loss of access to this key data source may go unnoticed by the risk analysts and result in an inaccurate score.

The final benefit of continuous analysis is risk pattern identification. This can be accomplished by analyzing the overall risk scores and sub-scores presented over a long period and establishing a general baseline. An example of risk pattern analysis would be identifying a potential organized botnet campaign targeting a key network based on deviations of the number of observed risk factors associated with incoming botnet DNS lookups. A significant uptick in a short period of time may indicate that the network in question is being specifically targeted and its risk for subsequent exploitation is higher than normal.

RISK CATEGORIES

Risk factors are binned into separate risk categories based on the type of risk to the network they present. Each category may also be assigned a separate severity level and score impact weighting based on its inherent risk to the network and the needs of the network owner. An example of potential risk categories is as follows:

Risk Category	Example Risk Factors	Severity Weighting
Software or Device Vulnerability	• Vulnerability to Windows SMB Exploits • High CVSS Scoring Software CVE Vulnerability	High to Critical
IT Hygiene	• Open and Publicly Accessible Remote Administration Ports • Non-Encrypted FTP Services	Low to Medium
Adversarial Threat	• Observed DNS Resolutions Toward Owned Domains by Identified Botnet IPs • Darknet Indications of Owned Email Credentials for Sale	Medium to High
Malicious Activity	• Owned Network Assets Communicating with Identified Malicious Command and Control Infrastructure	Critical

MARKESMAN ADVANTAGE

Markesman Group can leverage the insights and experience gained through its many experience cyber threat and risk analysts to tailor the perfect automated cyber risk analysis platform to any government or commercial customer's needs. The threat landscape and the risks associated with it are constantly changing. By codifying our cybersecurity expertise into your risk analysis platform we will be able to integrate the valuable experience we have gained from years of investigating and remediating threats to our nation's most sensitive networks.

Leveraging vRealize Automation in the Enterprise IT Operations Realm

What is vRA?

vRealize Automation or "vRA" is a portion of VMware's suite of tools, specifically focused on automation and orchestration of IT Assets, within pre-production, production, on-prem, and/or Cloud hosted services. In today's world of increasingly decreased deployment times and increased deployment frequency, especially those boasted by cloud companies such as Google and Amazon Web Services (AWS), automation integration in your environment is paramount to remaining viable in an ever increasingly complex and fast-paced IT world.

Markesman Group is capable of integrating vRA into your environment as a service provider in consideration of service delivery success and long-term sustainability. Markesman Group has deployed vRA enterprise wide for the Department of State and has the know how to effectively plug the tool into the environment, while bridging the gap between all stakeholders including VMware, internal customers, and asset holders.

Case in Point

In the IT Federal Space, it's not uncommon to experience weeks, sometimes month's long, lead and process time for basic IT Service Requests. Oftentimes, the request processes take far longer than the actual technical execution of the service being requested (ie: A basic request for a Windows Virtual Machine via Remedy Request Console). In this instance, it's quite common for the request queue to traverse multiple Government approvers, most of whom would otherwise approve the request (save for anomaly requests for retired Operating Systems and/or larger than necessary amounts of storage). These requests are largely performed manually, and as we all know, will sit in someone's approval queue going unnoticed, thereby extending delivery times. Fortunately, vRA integration with Remedy (or REST API calls to ServiceNow for customers hosting ticketing systems in the Cloud) can help to drastically cut down approval wait times, eliminating the need for human interventions. This allows for all approving stakeholders to focus on other strategic-type tasks, verses spending hours buried in a manual request queue.

Understanding Native Core Service Integration

For many clients, Windows and RedHat core services reign supreme as the heartbeat of their IT environment. Core services, such as SCCM and Windows Active Directory for Windows patch and configuration management, DNS, and Directory/Domain Service in conjunction with tools such as Satellite, Puppet, and RedHat IDM on the Linux side, are essential in keeping environments operational.

Though, as essential as these services are to VM provisioning, they are oftentimes executed through segmented teams operating in silos (ie: one team manages the Windows core services, another team spinning up a vanilla VM via master copy image, another checking Active Directory for Hostname availability, another registering the host with a DNS record, another installing third party tools, and another installing SQL, and another joining the server to a domain, etc.) Just to name a couple.

On the Linux side, IT Enterprise Teams may have a similar setup, with one team spinning up a vanilla VM via master copy image, and another registering the VM with Satellite and Redhat IDM domain services. These examples highlight work that may have automated components (such as vCenter executing a VM build via master copy image of "Template"), however the process to deliver a fully provisioned VM may pass multiple teams. This introduces human error while drastically increasing lead time between tasks, especially if engineers have developed email alert fatigue.

Introducing vRA into your environment reduces manual intervention and can eliminate this issue.

As an example, picture an on-prem environment that's distributed between multiple vSphere clusters and hosted in multiple network enclaves. From an engineering and administrative perspective, if this environment operates in silos as mentioned above, delivery of a basic Windows or Linux Virtual Machine may take several weeks.

To further highlight this example, let's assume there is a vSphere Cluster 1 and vSphere Cluster 2 in the same site, running in different networks, but with uplinks to the same Layer 3 Appliance. The customer has also requested that 100 Windows Virtual Machines to be provisioned on each Cluster within a week. Fortunately, with vRA implemented in your environment and integrated with your core services, this is more than possible. In fact, depending on your vCenter's bandwidth, this entirely likely within 1 business day. In this scenario, let's assume vRA picks up the task the moment the ticket is opened. From here the VM provisioning, registration, and integration with core services is automated using VMware's in-house tool called "Guest Script Manager" to execute automated commands and scripts, but via the ESXi/vCenter cluster hosting the VM to be provisioned.

Once provisioned, vRA will perform further tasks remotely to check for hostname availability, ensure the provisioned server is added to the collections group in SCCM, correct OU in AD, necessary security configurations applied, registered with DNS, and joined to the domain. The workflow and logic behind setting this up is more involved

than entailed; however, the end result is a fully provisioned VM on the network that's fully hardened, manageable, remotely accessible, added to patch/configuration management schedules, group policies with necessary 3rd party tools and application installed. If the customer requires this task to be performed 100 times, vRA not only introduces automation, but also consistency to each build, eliminating the need for human intervention.

Working with Markesman Group

The inherent benefits of leveraging vRA in your enterprise IT environment are nearly endless. While the goal of any vRA implementation is expediency, lessening the need for manual intervention, and faster time to market, Markesman Group understands and respects the tremendous complexity with end-to-end deployment of an automation tool in an enterprise IT environment.

As far reaching as these benefits are, enterprise-wise implementations are extremely complex. While Markesman Group is not intimately involved in your environment, we bring substantial experience/knowledge with managing/deploying an enterprise-wide vRA solution. This experience is vital, especially to identify potential blockers to success.

- Are your environments consistent across the board? Mismatched environments, such as those leveraging different directory service solutions, introduce complexity and increased logic behind the automated workflows. Markesman Group will apply strong Solutions Architecture experience to forefront to helps ensure consistency/limit unique "one-off" environments.

- Are your current native toolsets compatible with vRA If your native IPAM solution is not compatible with vRA, Markesman Group will help bridge the gap and implement vRA's internal IPAM solution, such as the need to create internal Network Profiles and the inherent risk that comes with that.

- Are you willing to spend the upfront capital needed to limit technical debt down the road? We all understand the importance of meeting deadlines and Markesman Group understands short term decisions have long term consequences. We will be transparent if the approach is simply "putting on a band-aid" as a temporary solution that impedes long-term sustainability and future automation enhancements.

- Documentation of Repeatable Processes. Realistically, customers cannot keep Professional Service vendors, like VMware, on their payroll over the long term. Therefore, it's essential to document any and all repeatable processes that will be necessary down the road, and can be executed by Tier 1 staff.

Lastly, how can we assist?

Inserting software and logic into currently manually executed workflows is anything but simple. That said, Markesman Group can help cut through these complexities, bringing vRA-Subject Matter Expertise, Systems & Solutions Architecture, Systems Administration & Engineering, and Project Management support to initiate, implement, monitor and continuously maintain the solution. We'll take the upfront time to address the bullets mentioned above, not only gaining a deep understanding of your environment, but a clear/upfront assessment of what it will take to achieve automation through vRA.

If you're interested or want to know more, please contact us.

Alex Wang	*Chief Operating Officer & Co-Founder*	Alex.Wang@Markesman.com
Charlene Polk	*Contracts Manager*	Charlene.Polk@Markesman.com
Timothy Hopkins	*IT Project Manager*	Tim.Hopkins@Markesman.com
Alex Drummer	*Cyber Lead*	Alex.Drummer@Markesman.com

GSA IT70 Schedule

Markesman Group offers cutting-edge cyber, information technology, cloud & digital transformation, systems development, customer programming, and intelligence services through the Markesman GSA IT Schedule with from various categories including: Program Managers, Software Engineers, Systems Analysts, Cybersecurity Specialists, Network Engineers, and more.

Click here to visit our GSA Schedule: Contract# 47QTCA18D00G0