Assessing the Operational System Risk Imposed by the Infrastructure Deployment Pipeline Workflow

By: Steven W. Porter, Raytheon IIS

Cyber Security & Information Systems Information Analysis Center

3 November 2020

TABLE OF CONTENTS

ABSTRACT

Real-time data monitoring of systems and system forensics is an essential aspect to keeping your Data Security Platform safe when relying on the use of Infrastructure as Code (IaC) and the potential vulnerabilities associated with its Continuous Deployment (CD). Many organizations are facing an information overload and are inadequately prepared for understanding and designing a cyber incident response plan with near-real-time monitoring, to include detection, analysis of system event logs, user activities and system access tracking.

A generalized Infrastructure Deployment Pipeline (IDP) reference architecture is presented to assist with risk assessment and mitigation. An experiment was conducted to determine if application of the National Institute of Standards and Technology (NIST) Cyber Security Framework (CSF) can mitigate the risks inherent to the IDP workflow process. The author concludes that while the NIST CSF does largely mitigate IDP cybersecurity risks, additional controls are still required to fully assure cybersecurity for the CD process.

This document also describes the benefits of Infrastructure as Code, and how to leverage the capabilities in support of DevOps (combined Development and Operations) initiatives. Infrastructure as Code is an emerging and evolving concept for automating the provisioning of infrastructure services and for managing infrastructure platforms such as virtual machines, networks, load balancers, and connection topology. The practice of Infrastructure as Code could be used as a catalyst/tools to increase organizations' abilities to deliver applications and services at a high velocity.

Additional guidance is provided for development teams to accelerate processes to enable rapid code production and deployment; and to assist in developing a vigorous agile strategy geared to deliver secure capabilities faster when relying on the infrastructure deployment pipeline (IDP) process. Moreover, this report describes several research studies that have addressed cybersecurity topics relating to IaC and IDP; and it details risk statements with architectural relationships of typical code signing solutions. Finally, it provides references on cloud computing services and scalable infrastructure resources.

INTRODUCTION

Imagine that the Stuxnet malware was not directly introduced to Iran's information systems as malware. Rather it was targeted at a United States information system and injected into a continuous deployment source code repository.

Once in the source code repository, the malware was built and deployed using elevated system privileges as part of a secure infrastructure deployment pipeline. With the widespread use of cloud computing services and the reliance on highly scalable infrastructure resources, the use of Infrastructure as Code (IaC) has become a requirement. Many of the opensource and even commercial software products require direct internet downloads in order to apply updates and patches. The use of continuous integration (CI) and continuous deployment (CD) has become standard practice for Information Technology (IT) organizations across all industries. The Infrastructure Deployment Pipeline (IDP) is a special case of the CD process that automates the provisioning of information system (IS) resources and enables rapid changes to the operational configuration in a consistent and reliable manner. The IDP workflow often requires elevated privileges to execute. The IDP can be a threat vector to the operational system. Even when the underlying IS are secured, the IDP workflow scripts and artifacts can be insecure. Combining an insecure workflow with the need to execute IDP processes with privileged user authorizations creates a threat vector that is not well understood. The National Institute of Standards and Technology (NIST, 2013) publishes national standards for cybersecurity which should be applicable to all types of IS systems. The NIST Cybersecurity Framework (CSF) is the primary cybersecurity roadmap for the United States Government and has been adopted by many non-government organizations. The IDP workflow is a security vulnerability to the operational system that can be mitigated by the application of the NIST CSF.

PREVIOUS RESEARCH

There are several research studies that have addressed cybersecurity topics relating to CI/CD, IaC, and IDP.

The article, Where are the Gaps? A Systematic Mapping of Study of Infrastructure as Code Research (Rahman, Mahdavi-Hezaveh, & Williams, 2019) was the impetus for IDP cybersecurity research. This paper, published in December 2018, identified the lack of academic research in the area of IaC cybersecurity practices and standards.

The article Security Support in Continuous Deployment Pipeline (Ali Babar, Zahedi, Ullah, Shahin, & Raft, 2017) focused on the specific topic of CD pipeline security. The paper addressed best practices for application continuous deployment but did not specifically address the infrastructure deployment aspects of CD. The research highlights the impact of application security risks in the Continuous Deployment Pipeline (CDP) and the risks of developers being able to subvert the pipeline deployment controls because of broad system privileges.

The article Securing a Deployment Pipeline (Bass, Holz, Rimba, Tran, & Zhu, 2015) took a Systems Engineering approach to securing the CDP. The article presented CDP requirements and modeled the pipeline using a secure supply chain methodology. This article did not specifically address infrastructure deployment requirements. This was one of the first research papers to present a CDP cybersecurity threat model. The model primarily addresses pipeline hardening requirements based on the threat of a remote attacker.

The Framework for Improving Critical Infrastructure Cybersecurity (NIST, 2012) provided a hardening framework that could be used to address the specific cybersecurity vulnerabilities posed by an IDP workflow on the operational system. The framework approaches cybersecurity as a risk analysis process for desired outcomes. This provides a high-level methodology to study IDP threat vectors across the complete SDLC. These four works form the basis for this research.

PROBLEM STATEMENT

Historically, IT operations System Administrators require special training and a higher level of trust because they possess elevated system privileges. By using current CDP methods, the insider threat risks have been transferred from IT operations to software development. The cybersecurity threat is now being compiled into the software on the very systems that are being protected.

The research presented in Where are the Gaps? A Systematic Mapping of Study of Infrastructure as Code Research (Rahman et al., 2019) provides a roadmap to new topics of research regarding IaC. One of the IaC publications reviewed by Rahman addressed IaC defects and security flaws (Rahman et al., 2019). There is little research available addressing IaC security or specifically the CDP workflow threat. The case study documented in A Deployment Pipeline for Infrastructure: A DevOps Case Study at NBN (Humble, 2009) presented an IDP use-case in the context of the Software Development Life Cycle (SDLC). Including multiple development teams, unit testing, automated code promotion, configuration specifications, and deployment smoke testing. The case study raised many questions about cybersecurity applications within the deployment pipeline. It was obvious that the supporting system infrastructure for the pipeline required cybersecurity hardening, and that the operational system infrastructure required cybersecurity hardening. There appears to be an area between the deployment pipeline and the operational system that represented a gap in current research. If the deployment pipeline workflow processes were left unconstrained, then a vulnerability introduced in the infrastructure deployment pipeline workflow could escape into the operational system. The infrastructure threat impact could be considerably higher in the operational environment than for an application because most infrastructure deployment processes require elevated system privileges.

This research addresses two primary questions. First, what is the IDP workflow risk created by an identified vulnerability? Second, does the IDP workflow represent an insider threat to the operational system?

RESEARCH METHODOLOGY AND DESIGN

A generalized infrastructure deployment pipeline reference architecture provides context for the risk assessment. The research data collection is constrained by the IDP system reference architecture, the age of the data, and data source.

There is not an industry-standard IDP reference architecture; however, there is considerable academic agreement on the stages of the deployment pipeline. This reference architecture was chosen because it represents the operational system viewpoint of the IDP workflow and is based on previously reviewed academic research. (Ali Babar et al. 2017) Source code produced by the developer is committed to a branch in the source code repository. The successful commit of source code triggers the CI service to build the appropriate binary artifact. The CI service allocates the build to the appropriate Build Service. Once the appropriate artifact is built, the CI service sends a request to the Test Service. The Test Service executes a series of tests depending on the level of testing required. Following successful testing, the artifact is ready to be deployed (Ali Babar et al.). In addition to the Babar reference architecture, a binary repository and continuous deployment service were added to represent an IDP lifecycle, as shown in Figure 1.

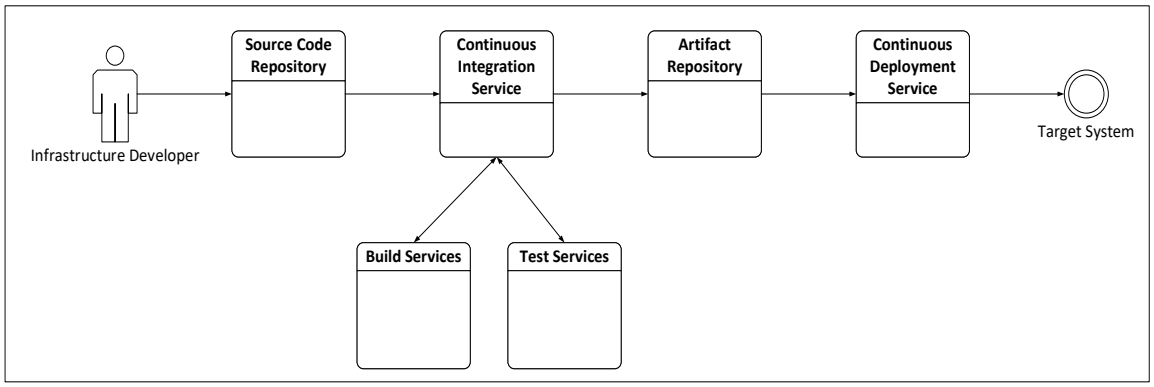

Figure 1: Infrastructure Deployment Pipeline Reference Architecture

Infrastructure Deployment Pipeline Tools

There are many tools that fulfill the infrastructure deployment pipeline capabilities. The DevOps community documents many opinions of best-practice tools, but there is not a definitive, scholarly reference. In a research paper Security Support in Continuous Deployment Pipeline (Ali Babar et al., 2017), a notional set of tools was defined that fulfilled each of the reference model capabilities. The tool vulnerabilities selected for research fulfilled the requirements of each component in the reference architecture (Ali Babar et al., 2017).

Cybersecurity Framework Reference Architecture

The NIST CSF was selected as the authoritative security reference framework for research. The CSF prescribes controls through information references. The information references specify sections of standards, guidelines, and practices documented in other NIST publications (NIST, 2012). For this research, the NIST SP800-53-R4 information references were used. The NIST SP800-53-R4 controls map directly to the CSF core requirements.

Sample Selection Criteria

The research population is the group of all cybersecurity vulnerabilities. A cybersecurity vulnerability is a defect in an information system that leaves the system open to an attack vector. To be authoritatively recognized as a member of this population, the Common Vulnerability and Exposure (CVE) database maintained by Mitre Corporation was used (Mitre, 2019). The CVE database is queried using the name of each IDP workflow tool identified in the IDP reference model. In order to maintain current-day risk relevance, only vulnerabilities that are less than five years old where selected. The vulnerability must be relevant to the IDP workflow or deployable artifacts. For each vulnerability matching the selection criteria, the researcher composed a risk statement. The risk statement ensured that only the portion of the vulnerability that applied to the IDP workflow was used in the research.

In order to assess the magnitude of the risk for each of the selected CVE records, the National Vulnerability Database (NVD) was queried using the CVE number. The NVD provides enhanced information for each valid CVE vulnerability, including fix information, severity scores, and impact ratings. The NVD also categorizes the vulnerabilities by vendor name, application, operating system, and threat vector. The NVD data is used to populate the risk register with the user privilege value, system access method value, confidentiality, integrity, and availability threat values from the NVD database.

Threat Data Analysis Methods

The identified vulnerabilities are grouped according to threat vectors. Using threat vectors provides a construct that is appropriate for categorizing vulnerabilities. Vulnerabilities tend to be very technical and specific to the state of the system. Using the threat vector abstraction allows the vulnerabilities to be grouped by themes. Each identified threat vector includes multiple vulnerabilities with an associated risk.

The Confidentiality, Integrity, and Availability (CIA) impact on the IDP workflow and operational system is analyzed for each identified risk. The threat impact is defined by the Exposure Factor (EF) and is stated as a percentage of change in the operational CIA security posture. The EF provides a construct for quantifying the impact of the threat vector on the operational system. The operational system is abstracted to describe just the delta in CIA posture caused by the impa ct of a vulnerability. The EF value is a subjective percentage of the operational asset lost due to the impact of the vulnerability (Karabacak & Sogukpinar, 2005). The EF is used to calculate the Single-Loss Expectancy (SLE) value. The SLE is calculated as the Asset Value (AV) multiplied by the EF. For this research, the AV will always have a value of one hundred (Karabacak & Sogukpinar, 2005). The EF is used to assess the impact of a realized IDP workflow risk on the operational system.

The research analysis answers two questions.

a) How does the IDP workflow risk compare to the operational system risk for each vulnerability?
b) Does the application of the CSF mitigate the insider threat to both the IDP workflow and the operational system?

Threat Vectors

The vulnerabilities and associated risks were mapped to common cybersecurity threat vectors. The threat vectors, as defined by the Carnegie Mellon University Software Engineering Institute, were used as the basis for classification (SEI, 2015). The threat vectors for the identified IDP workflow vulnerabilities are malicious code execution, unauthorized data access, weak passwords, elevated user privileges, compromised credentials, and denial of service. The distribution of vulnerabilities across the various threat vectors is shown in Table 1.

Table 1: Threat vector mapping of vulnerabilities

Threat Vector	Count
Malicious Code Execution	7
Unauthorized Data Access	10
Weak Passwords	1
Elevated Privileges	6
Credentials Compromised	6
Denial of Service	1

The results of the threat vector analysis show that the vulnerabilities tend to be grouped around a small number of threat vectors. The data sample provides multiple examples of each type of threat vector. The threat vectors are common to both the IDP workflows and the operational system.

Operational System Exposure Factor

The EF is the subjective delta in the operational CIA posture caused by the impact of a vulnerability. The EF is assigned based on the impact of the IDP workflow threat on the operational system. The EF was determined using the guidelines documented by SANS for estimating EF and emphasizing the impact component of the IDP workflow risk (Tan, 2003). The impact component of the calculation was emphasized because the likelihood component relies on the state of the IDP for initial vulnerability realization. The likelihood remains constant across the SDLC of the system. The initial analysis of the EF values shows that the IDP vulnerabilities tend to be high impact threats for the operational system.

Once the EF for each vulnerability has been calculated; the SLE risk value can be calculated. The value of the operational asset is fixed to a value of 100 so that the impact of the EF independent variable can be examined. The Single-Loss Expectancy value for the operational system was calculated for each of the identified vulnerabilities.

IDP Workflow Risk

A risk statement was written for each IDP workflow vulnerability identified. The risk statement ensured that only the part of the vulnerability associated with the IDP workflow was considered in the analysis. For each risk, a risk magnitude value was calculated. The risk magnitude value is the product of the vulnerability Likelihood and the

vulnerability Impact values. The risk magnitude value quantifies the level of risk each vulnerability has on the IDP workflow. The IDP workflow risks tend to be high impact threats.

IDP Workflow Risk Magnitude and Operational SLE Comparison

The risk magnitude value and SLE are both measures of system risk posed by the vulnerability (SEI, 2015). Comparisons can be made between the risk of the vulnerability to the IDP workflow and to the operational system.

In order to make a relative comparison between the IDP workflow risk and SLE risk, the values must be standardized by calculating each values Z-score. By using Z-scores, risk values can be compared using the same scale. When the IDP workflow z-values are plotted with the SLE z-values, the pattern indicates independence between the IDP workflow risk impact and the SLE risk impact, as shown in Figure 2.

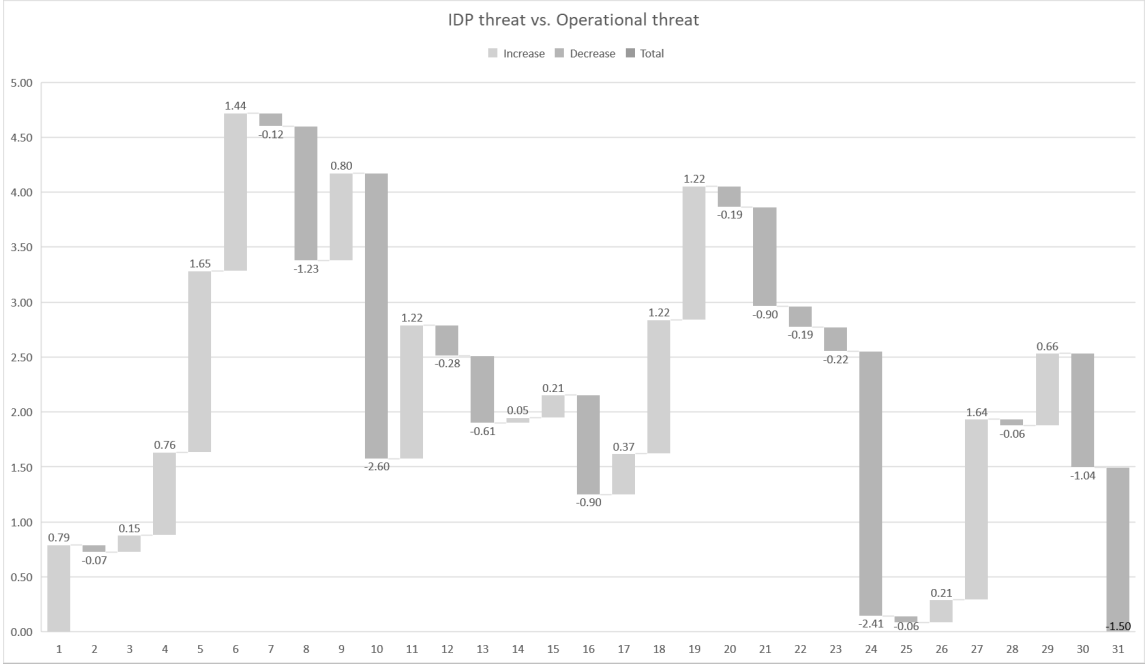

Figure 2: IDP Risk z-value vs. SLE Risk z-value

The graph indicates the relative difference between the IDP risk and the SLE risk. The positive values indicate that the risk impact is greater for the IDP workflow than the operational system. The negative values indicate that the risk impact is greater for the operational system than for the IDP workflow. These are IDP workflow vulnerabilities that pose the greatest threat to the operational system if they are allowed to escape. The longer the graphs vertical bar is, the greater the difference in risk impacts.

An analysis of the vulnerabilities indicates that the vulnerability likelihood remained the same for the IDP workflow and the SLE. This is because the complexity and privilege requirements remain the same independent of where the attack is targeted, but the impact of the attack is different depending on the asset target.

The comparison answers the question of differences between the IDP workflow and operational system impact. Though there is not a statistical correlation, there is an interdependence between the risks. The IDP workflow vulnerability must be realized for the SLE attack vector to be realized. The realization of the IDP workflow vulnerability does not necessitate a vulnerability realization on the operational system.

IDP Workflow Risk Mitigation using CSF

Based on the cybersecurity vulnerability and associated risk definition, each data set element was mapped to a set of CSF Functions and associated Categories which define the scope of the risk mitigation outcomes. The research applied the NIST SP800-53-R4 information reference as documented in the NIST Special Publication (SP) 800-53 Revision 4 (NIST, 2013) to determine the applicable controls to apply for risk mitigation of each risk category identified.

The application of the cybersecurity controls does mitigate the IDP workflow cybersecurity risks. The application of the cybersecurity controls requires several key architectural patterns for mitigating the threat vectors.

a) To prevent the introduction of malicious code that could threaten the operational system, all deployable artifacts must be built from a version-controlled source code repository.
b) The IDP must require multiple levels of access. Development, test, and operational deployment should require different RBAC permissions, and no person or process should have access to all of the roles.
c) Authentication and authorization, including strong passwords and certificate management, should be managed at the enterprise level with consistent policy and training enforced across development, test, and production.
d) Elevated privileges must be managed at the enterprise level. Each role must have the minimum privileges to execute their job. No person or process should have full access.
e) The IDP should be isolated within the operational system to prevent Denial of service attacks.

The identified threat vectors present an asymmetrical threat to the operational system. The analysis indicates that in many cases, a low-risk to the IDP workflow may pose a much higher risk to the operational system. There is at least an architectural gap in the application of cybersecurity controls. The CSF does not directly mitigate the impact of defective scripts promoted by the IDP to the operational system. There are three main defect types that define defective scripts which are filesystem operations, infrastructure provisioning, and user account management for an IDP workflow (Levet, Granier, & Schlick, 2006). The categories and controls lack the ability to validate that the IDP only executed the proper operations. There is also inadequate assurance that the automated management of user accounts only provisioned the correct accounts. This is an unmitigated threat vector because most IDP scripts require elevated privileges.

CONCLUSION

The research results conclude that the application of the CSF controls does largely mitigate the IDP workflow cybersecurity risks, but a partially unmitigated threat vector remains for the operational system.

Many of the IDP workflow vulnerabilities present an asymmetrical threat to the operational system. After the CSF categories are evaluated, there is a substantial IDP workflow risk to the operational system due to the exposure to defective IDP scripts that execute with elevated system privileges. The CSF and NIST SP800-53 controls must be evaluated specifically for each the IDP, the IDP workflow, and the operational system even when the IDP is within the operational system security boundary.

Additional cybersecurity controls may be required to adequately address the cybersecurity threats posed by the IDP workflow. These include additional separation of duties and least privilege access controls (NIST, 2013). The objective is to provide an approved development process, a predefined set of tools with real-time monitoring designed to achieve a high availability by minimizing time to detect and time to mitigate to the cybersecurity professional teams via automated monitoring.

Secure coding standards specific to infrastructure code need to be developed. The research revealed secure coding standards for C, C++, Java, and Perl published by SEI. Research into secure coding standards for the primary Domain-Specific Languages (DSL) such as Chef DSL and Terraform DSL did not reveal any authoritative guidance.

REFERENCES

Ali Babar, M., Zahedi, M., Ullah, F., Shahin, M., & Raft, A. J. (2017). Security support in continuous deployment pipeline, 57–68. https://doi.org/10.5220/0006318200570068

Bass, L., Holz, R., Rimba, P., Tran, A. B., & Zhu, L. (2015). Securing a deployment pipeline. Proceedings - 3rd International Workshop on Release Engineering, RELENG 2015, 4–7. https://doi.org/10.1109/RELENG.2015.11

Humble, J. (2009). A deployment pipeline for infrastructure: A DevOps case study at NBN, (December), 1–7.

Karabacak, B., & Sogukpinar, I. (2005). ISRAM: Information security risk analysis method. Computers and Security, 24(2), 147–159. https://doi.org/10.1016/j.cose.2004.07.004

Levet, F., Granier, X., & Schlick, C. (2006). Anti-patterns in infrastructure as code. Lecture Notes in Computer Science, 4073, 114–125.

Mitre. (2019). Common Vulnerability and exposures. Retrieved from https://nvd.nist.gov/

NIST. (2012). Framework for improving critical infrastructure cybersecurity. Proceedings of the IEEE, 100(1), 210–224. https://doi.org/10.1109/JPROC.2011.2165269

NIST. (2013). Security and privacy controls for federal information systems and organizations, 4. Retrieved from https://csrc.nist.gov/csrc/media/publications/sp/800-53/rev-4/archive/2013-04-30/documents/sp800-53-rev4-ipd.pdf

Rahman, A., Mahdavi-Hezaveh, R., & Williams, L. (2019). A systematic mapping study of infrastructure as code research. Information and Software Technology, 108(i), 65–77. https://doi.org/10.1016/j.infsof.2018.12.004

SEI. (2015). Threats and risk calculations, 1–22.

Tan, D. (2003). Qualitative risk analysis step-by-step. SANS.

ABOUT THE AUTHOR

Steven W. Porter

Steve Porter is an Engineering Fellow at Raytheon IIS. Steve is a subject matter expert on DevOps, Infrastructure as Code, Cloud Computing, and infrastructure architectures. He has worked at Raytheon for twenty-four years supporting a variety of domestic and international programs. Steve has a Master of Science in Cyber Security Management degree from Purdue University Global and Bachelor of Science in Engineering Technology degree from the University of Northern Iowa. Steve is a Raytheon Certified Architect.

PREPARED BY:

Cyber Security & Information Systems
Information Analysis Center

Operated by Quanterion Solutions Incorporated
under contract FA8075-17-D-0001

266 Genesee Street
Utica, NY 13502
https://www.quanterion.com
qinfo@quanterion.com