REPORT

WHY EVERY FEDERAL GOVERNMENT ORGANIZATION SHOULD HAVE AN ENTERPRISE APPLICATION STRATEGY

March 2019

CONTENTS

INTRODUCTION

Federal IT executives must walk a fine line to deliver more value while still reducing waste and minimizing losses due to cyber-attacks. On the one hand, IT leaders must embrace the latest technology to unlock greater efficiencies, deliver more value, and save taxpayer money. And, on the other hand, they face an evolving and expanding threat landscape requiring significant resources to mitigate. In short, with every initiative and every decision, they must both innovate and protect. To navigate these seemingly opposing goals, Federal IT leaders should develop an enterprise application strategy, tailored to their organizational mission.

DIGITAL TRANSFORMATION IS USHERING IN THE AGE OF APPLICATIONS

Digital transformation is a new way for organizations to use technology to enable profound transformations in performance. In short, applications are becoming the business or mission itself. To take advantage of this radical shift, organizations are making sweeping changes in pursuit of new business and IT operational models. 65% of respondents surveyed in F5's 2019 State of Application Services research reported that their organizations are in the midst of a digital transformation.

Digital transformation is driving rapid growth in the number of applications, that is software designed to perform a specific function directly for the user, or to support another application. Applications power the growing utility of smart phones, they deliver business-critical functions like CRM, and ERP, and everything in between. IDC estimates there are 250M+ enterprise application workloads alive in the world today, growing to 1,700M in 2022. The 2018 F5 Labs Application Protection report revealed that the average public sector organization uses 680 applications, 32% of which are considered mission critical.

On top of this tremendous growth in applications, organizations are deploying those applications into new architectures (e.g., microservices, serverless) and into new environments (e.g., public cloud). F5's 2019 State of Application Services research indicates that 14% of organizations surveyed have made containers the default application workload isolation approach and 87% of respondents are adopting a multi-cloud architecture. The Office of Management and Budget (OMB) itself is advocating a "Cloud Smart Strategy" to migrate to a safe and secure cloud network.

RAPIDLY EXPANDING APPLICATION THREAT SURFACE AREA

But government and commercial organizations aren’t the only ones using technology to reduce friction, supply new services, and enhance value. Cyber-criminals, militant groups, and nation-state threat actors are also innovating their cyber-attack capabilities at breakneck speed. As we have gotten better at protecting our networks and infrastructure, cyber-attackers are shifting their sites to softer targets. F5 Labs research shows that 86% of all cyber-attackers either target applications directly or steal user identities, usually by phishing. In fact, spear phishing privileged users is how nation-state APTs in Russia and China typically begin their attacks on targets of interest in the US.

Attackers are learning that it’s sometimes easier to attack indirectly. One way is by going after less-important applications, such as IoT devices, and leveling up their access. Such was the case where cyber-crooks hacked a US casino through a lobby aquarium thermometer and used it as a foothold into the high roller database. The devastating Target breach of 2013 began as an exploit of Target’s HVAC vendor and ended with a breach of the point of sale system and a loss of 40 million credit cards.

Nation-state attackers highly proficient at targeting have moved upstream to strike at third-parties and supply-chain providers of the US government. It has been reported that the 2015 breaches at Anthem and the Office of Personnel Management (OPM) were part of a larger initiative to obtain data on the CIA. The attackers sought to do a diff on the data collected from Anthem, which included all government employees, and OPM, which did not have data on CIA operatives. Such sophisticated and well-planned attacks pose a real and increasing threat to national security.

Left unmanaged, Open Source and other third-party components or services can also introduce risk. The Sonatype State of the Software Supply Chain 2018 notes that 1 in 8 Open Source component downloads contain a known security vulnerability. Organizations, including in the public sector, take advantage of Open Source because it speeds up development and delivery. However, these same organizations often fail to include Open Source components in security scans and review processes. These components become part of the application portfolio and should be treated with the same scrutiny as custom code. This includes components on both the “server” side (e.g., NPM packages, libraries) as well as the “client” side.

With applications now playing a more critical role in fulfilling the goals of every organization, the growth of those applications out-pace the ability for most organizations to scale their operations. New architectures and deployment models further challenge traditional security measures. Meanwhile, the threat surface area associated with applications is expanding rapidly.

MULTI-CLOUD APPLICATION SERVICES ARE JUST ONE PART OF THE SOLUTION

As if things weren’t complex enough, organizations are moving to multi-cloud deployments to leverage greater flexibility, gain higher availability, reduce vendor lock-in, and, in some cases, take advantage of the pricing arbitrage opportunities between clouds. The quandary is: how to deliver a standardized, secure, and seamless application experience across hybrid and multi-cloud architectures without exceeding your budget allocation.

Standardizing on multi-cloud application services reduces the operational complexity that comes along with a multi-cloud architecture. For example, it's easier to manage and more powerful to have a service that functions at an enterprise-application level, such as code-based invocation of APIs and event-driven systems, than a service that is platform-specific, such as a having to use a vendor-specific message queuing service. Enabling one set of features across all applications cuts down on operational overhead. By adopting a standard platform for as many application services as possible, organizations can leverage more automation and code reuse to realize consistent, predictable, and repeatable operational processes.

DEVELOPING AN ENTERPRISE APPLICATION STRATEGY

In an age in which many born-in-the-cloud businesses don’t even have an IT department, traditional IT architectures and operational processes fall noticeably short of application developers and DevOps teams’ expectations. The desire to drive applications out faster often leads to bypassing traditional NetOps and SecOps teams along with the associated security and operational processes. Indeed, protecting the enterprise application portfolio has as much to do with people and processes as it does with technology.

To manage performance and, more importantly, risk across this multi-cloud sprawl, organizations are desperate for solutions. Wherever the applications live, solutions must support the deployment of consistent policy, manage threats, and provide visibility, and allow for monitoring of health and performance. A poor-fitting solution could easily diminish any benefits from Digital Transformation if the innovation and agility of application development and DevOps teams is hindered. Given the increasing value (and risk) profile of applications, every organization should develop an Enterprise Application Strategy that addresses how applications in the enterprise portfolio are built/acquired, deployed, managed, and secured.

Step 0: Align Application Strategy goals with the organization mission

To be effective, the Enterprise Application Strategy must be in support of the organizational mission. After all, the whole point of Digital Transformation is to replace unwieldy or manual processes with hyper-charged, data-rich applications. Therefore, the overarching goal of an Enterprise Application Strategy should be to directly enhance, speed-up, and protect the organization's digital capabilities as they relate to delivering against the mission. Any application or associated application service that is incongruent to this goal should be deprioritized.

This alignment also means taking the status-quo into consideration, which includes looking carefully at the current enterprise data strategy, compliance requirements, and the overall risk profile of the organization. In many cases, the set of constraints imposed by these different sources and the impact on the innovation and agility of application development teams and their compliance with those constraints is likely not well-understood. The Enterprise Application Strategy should clarify the balance the organization is willing to strike between the often-competing forces of innovation, agility, and risk.

Step 1: Build an application inventory

When it comes to an Enterprise Application Strategy, most leaders are never lucky enough to start afresh. Nearly everyone in the IT industry inherits a technology architecture that is a result of decades of disparate systems mashed on top of legacy systems coaxed to keep functioning. This problem can be especially severe within the Federal government, as compared to the commercial sector. It is rarely easy to cleanly migrate these accretions of technology to a desired target state. Therefore, more discovery and analysis must be done.

It may sound silly, but to adequately protect something you must first know that it exists, and then be able to accurately monitor its health. And yet, with few exceptions, organizations are not able to report with confidence the number of applications they have in their portfolio, let alone if those applications are healthy and secure. The F5 Labs 2018 Application Protection Report discovered that 62% of IT security leaders have low or no confidence in knowing about all the applications in their organization.

An enterprise application inventory is the most foundational element of any application strategy. This is a catalog of all enterprise applications, whether delivered internally, laterally (e.g., to other government entities), or externally (e.g., to the public), that includes:

- A description of the function the application or digital service performs
- The origin of the application (e.g., custom developed, packaged software or third-party service)
- The key data elements the application requires access to or manipulates

- Other services the application is communicating with
- Open Source and other third-party components that are part of the application
- The individual(s) or group(s) accountable for the application

Building the application inventory for the first time is often painstaking and time-consuming work. One approach for easily smoking out rogue applications is by making the application inventory a whitelist – applications that are not on the list simply do not get access to enterprise resources (e.g., network). To chase down applications outside of your organization, a tool like a cloud access security broker (CASB) can be very helpful. CASBs sit between your users and the Internet, monitoring and reporting on all application activity. They can not only tell you what the top applications your employees use (and how they access them) but also give insight into shadow IT application usage.

Every moment spent ensuring the accuracy of your application inventory has a direct impact on your ability to quickly define FedRAMP system boundaries. It also allows you to have a much more accurate way to quickly and precisely find an area Responsible Party for an application's infrastructure – or even an individual component of an application – when asked to do so by the certifying body.

Lastly, your FedRAMP efforts require that this is a continuous exercise. It may not even be enough to do it annually. It's a constant process that involves keeping an eye on what applications and data repositories are in play, monitoring what users need to do, and evaluating how your development environments are evolving.

Step 2: Assess cyber-risk for each application

Cyber-risk is a significant and growing concern for IT leaders. It is a deep and complex subject and it begins with evaluating the risk to the applications. Each of the applications in the inventory should be examined for four primary types of cyber-risk:

- Leakage of sensitive internal information (e.g., military secrets)
- Leakage of sensitive customer/user information (e.g., personnel records, tax history)
- Tampering of data or applications
- Denial of service to data or applications

In cyber-risk, the importance of a digital service is measured in the risk equation as the financial or reputational impact to that service from the cyber-attack per the categories discussed above. Different organizations will place differing levels of potential loss for some services than others, so each organization should do their own estimations based on their defined mission. FISMA, for example, asks you to determine agency-level risk to the mission or business case but going deeper to examine risks at an application level prepares you for when your mission's compliance standards inevitably deepen.

Sometimes included with organizational risk calculations are the risks of non-compliance with applicable rules, guidelines, and contracts. Federal entities are steeped in regulations and standards, and these should all be accounted for when evaluating applications and digital services. Establishing an application-linked model to assessing cyber-risk helps ease the process of satisfying the CDM/ConMon requirements to FedRAMP compliance by letting you narrow the boundaries to manageable and appropriately granular groups of services.

Step 3: Enumerate the application services needed

Applications rarely stand alone, so along with the application inventory, the application services should also be discovered and tracked. Application Services are packaged solutions for application builders that improve the speed, mobility, security, and operability of an application. These services bestow several important benefits to the application workload:

- Speed: The performance of an application workload and the ability to deliver quickly.
- Mobility: The easy movement of an application workload from one physical or logical hosting site to another.
- Security: The protection of the application workload and the data associated with it.
- Operability: The assurance that an application workload is easy to deploy, easy to keep running, and easy to troubleshoot if it fails.

Common application services include load balancing, DNS delivery, global server load balancing, web application firewall (WAF), DDoS prevention/protection, application monitoring and analytics, identity and access management, application authentication, API gateway, container ingress and egress control, and SSL encryption. All of these application services involve some level of cost, both direct in terms of the service itself and indirectly in terms of operational upkeep.

A good way to locate dependent application services is by examining the Controls section of the FISMA or FedRAMP System Security Plan for your environment. This will often point out the presence of both security-focused application services and other services that are dependent to them.

While every application can benefit from application services, not every application requires the very same application services. It is also worth considering that many application services are specifically designed to support narrow categories of applications. For example, only applications designed to serve IoT apps need an IoT gateway. Applications delivered in a traditional architecture don't require app services targeting containerized environments, so ingress control and service mesh application services may not be applicable.

In some cases, new application services may need to be acquired to ensure compliance or to reduce risk. Although your compliance regime allows for it, you should always resist the temptation to select the minimum baseline controls and insist on selecting application services that enhance your ability to deal with cyber-risk.

Step 4: Define application categories and assign minimum application service requirements for each category

When it comes to drafting the strategy, it is useful to map the stated goals of the policy to the applications, application services, and risks. Once the application inventory is complete, the next step is to group the applications into logical categories based on the characteristics that need different management and application service approaches (e.g., access to sensitive data, exposure to more threats).

Once categorized, the enterprise application policy should specify the performance, security and compliance profiles that are to be applied to different application types, based on the criticality and enterprise classification of the application itself. One common way of categorizing applications is as follows:

- Tier 1
  - Application characteristics: Mission Critical Digital Services that collect and transform sensitive data
  - Required application services: load balancing, global server load balancing, web application firewall, DDoS protection / prevention, bot detection, SSL encrypt and decrypt, user identity and access management, application / service identity and authentication, application visibility / monitoring
- Tier 2
  - Application characteristics: Mission Critical Digital Services that provide access to sensitive data
  - Required application services: load balancing, global server load balancing, web application firewall, DDoS protection / prevention, bot detection, SSL encrypt and decrypt, user identity and access management, application / service identity and authentication, application visibility / monitoring
- Tier 3
  - Application characteristics: Mission Critical Digital Services that do not collect or provide access to sensitive data
  - Required application services: load balancing, global server load balancing, DDoS protection / prevention, application visibility / monitoring

- Tier 4
  - Application characteristics: Other Digital Services
  - Required application services: load balancing, application visibility / monitoring

As the threats facing applications vary based on the environment in which they are hosted, this categorization can be further expanded to differentiate based on deployment environment (e.g., on-premises hosting versus public cloud).

Prioritizing your goals in this fashion also helps you “pre-classify” applications that you deploy into proper FISMA/FedRAMP levels easily. Spending a little time here developing a structure to your mission goals allows you to spend much less time talking to an auditor later.

No organization is going to have enough resources to do everything they want in an acceptable timeframe. With the prioritization of these applications, you can take a triage approach to what apps need bolstering with application services, which applications should be modernized or replaced, and which applications aren’t worth the effort. If the applications aren’t worth the effort, make sure they are segmented off in your network and avoid the scenario where an innocuous IoT thermostat leads to a full network breach. This process also includes looking at new applications that could unlock new value streams and should therefore be either developed internally or sourced from a third-party.

Step 5: Define parameters for application deployment and management

A foundational part of any IT strategy has always been deployment and operational management, and a modern Enterprise Application Strategy adds a few new twists (e.g., importance of the end-user experience). This includes looking at:

- Which deployment architecture(s) are supported (e.g., hybrid-cloud, multi-cloud)
- Deployment model options for each of the application categories
- Which public clouds can serve as access points for applications
- To what extent public cloud native services can be leveraged versus third-party

Finally, different applications have different needs in terms of deployment and consumption models. During this phase of developing your application strategy, you should strive to gain a clear understanding of the different deployment options, each of which might have different consumption models, cost impacts, and compliance / certification profiles. In selecting deployment models, it is also prudent to inventory available skills and talent to factor into the decision. For example, choosing to deploy onto AWS (or any cloud) when you have insufficient in-house talent to manage it and lack access to contract-based skills can slow you down and introduce risk.

Never forget that your deployment and management mechanisms may, themselves, be subject for an authorization, whether it be under FISMA or FedRAMP ATO/P-ATO, whichever your agency uses as a standard for your mission.

Step 6: Clarify roles and responsibilities

In addition to articulating the goals and priorities, the Enterprise Application Strategy should also include elements around roles and responsibilities. Key questions to answer here are:

- Who has decision rights around optimizing and securing the application portfolio (e.g., technology selection, application disposition, user access management)?
- Who has Privileged User Access to each application?
- Who will be responsible for deployment, operations, and upkeep of each application in the various environments?
- Who in the organization is responsible for compliance with the Enterprise Application Policy?
- Who is going to monitor for compliance to the Enterprise Application Strategy goals? And who will they report metrics to?
- Who is going to monitor vendors (including Open Source and third-party component / service providers) for compliance?
- Who is going to ensure all applications and application services are accounted for (as applications and services will keep changing and being added/removed)?

Some of these responsibilities may fall to entire departments while others may fall to multi-departmental committees or single roles. Whatever it is, these roles and responsibilities should be spelled out clearly; they may, in fact, need more definition than how your compliance regime requires you to define them. More advanced organizations will adopt operational processes and automation to assign these accountabilities early in the development process, at the time of application inception. In a multi-cloud world of hundreds or even thousands of applications supporting critical functions, the Application Strategy and corresponding policies should establish clear lines of accountability.

ENFORCING THE ENTERPRISE APPLICATION STRATEGY

Once the Enterprise Application Strategy is developed, to serve its purpose, it must be enforced. Enforcement mechanisms should include “hard” guardrails built into the automation of processes (e.g., user access control, code vulnerability scans at check-in) as well as “soft” measures such as employee training and capability- or awareness-building.

Implement robust access controls in support of the roles and responsibilities

Your access control policy should support the operational roles and responsibilities defined in the Enterprise Application Strategy and extend to all applications both on premise and in the cloud. Special attention should be applied towards privileged user access because of the risk they pose to the application, including being targeted by sophisticated APTs because of their administrative or root permission to the application. Special measures recommended for privileged users include:

- Privileged users should always be in separate groups, defined as “high risk", inside of your access control solution that require security controls you might not choose to implement for all users, or all application classification tiers.
- Multiple factors of remote authentication should be required. If access attempts are made with valid credentials that fail the second authentication requirement (in the case where attackers have collected valid credentials from a phishing attack or through another breach where credentials were shared), or from a location not physically possible based on last valid login (such as a successful login from a US office two hours prior to the same user attempting to log in from an Eastern European IP), the account should be locked until further security review is completed.
- Administrative access should only be authorized for appropriate, and trained, personnel that require this level of access on a regular basis to perform their job. Any temporary access granted for emergencies or special projects should be in a different user group set up with automated use monitoring that will remind system administrators if they forget to remove the access. Review of access appropriateness should be performed on a regular basis and completed independently of the team responsible for the application, or granting access to the application, to avoid any conflicts of interest. If a privileged user does not access an account for an extended period, it should be questioned whether they truly require the access.
- Proper accounting of all privileged user access to the applications should be logged, as well as any changes made by the user account.

Getting this level of visibility and automation around access control in the cloud can be challenging and costly as these features are generally not available natively. It is however possible with third party licensing and given the importance, its an investment well worth making.

Continuously train employees and relevant stakeholders

With the steady growth in applications, and the abundance of data available in the media that attackers use to figure out what applications to target and who has access to them, security awareness training has never been more important. As spear phishing is the modis operandi of adversaries, phishing training should be a large focus. The 2018 F5 Labs Phishing and Fraud Report found that training employees more than 10 times can reduce phishing success from 33% down to 13%. Yet rarely is security awareness training conducted enough, and with the right material. Canned awareness training services designed to check compliance boxes once a year run the risk of employees not understanding their role in information security, and not having a personal sense of duty to it. If the goal is to reduce the risk of a breach, frequent training, personalized to your organization, is the way to go.

Employees should be aware that cyber-attacks are a constant threat. There is no downtime for attackers, and thus employees must always remain vigilant. A continuous culture of curiosity adapted by the NSA should be the norm for all organizations, especially in the Federal space or any businesses supplying the Federal government with products and services. Employees should be aware that they are a target because of their access to applications and data. They should also be aware of how that access or data is used in advantageous ways by adversarial nation-states or sold by for-profit cyber criminals (that is then purchased by adversaries).

CONCLUSION

To ensure success in their digital transformations, all organizations should adopt an Enterprise Application Strategy and corresponding Policy, and train employees on them. Within the Federal space, with its large mixture of legacy, hybrid, and new-model applications, this is especially critical. Your success in delivering a reliable, secure, and authorized digital service requires it.

Applications are the heart of any organization’s digital transformation and, with the rapid change in the way software is developed and deployed, they are both an organization’s greatest source of value and the greatest source of vulnerability. The Application Strategy and Policy components outlined here provide the essential foundations to secure any organization’s digital aspirations. With the risk profile of their application portfolio only increasing with each day, organizations must move quickly to formalize their Strategy and Policy.