

A man with glasses and a blue lanyard with an ID badge is walking through a server room. He is wearing a dark sweater and dark pants. The server racks are black with blue and green lights. The background is a blurred view of the server room.

A PRACTICAL GUIDE TO REDUCING COMPLEXITY WITH AUTOMATION TECHNOLOGY

Table of Contents

| | |
|--|---|
| Modernizing the Data Center Through Automation | 3 |
| Benefits of Automation | 3 |
| A Mindset for Automation | 4 |
| Creating a Software-Oriented Team | 4 |
| Gathering Automation Requirements | 4 |
| Implementing Automation with the “Six Sevens” Best Practices | 5 |
| The Six: A Universal Multi-Cloud Data Orchestration Process | 5 |
| The Seven: Multi-Cloud Foundations | 6 |
| The VMware Advantage | 9 |

Modernizing the Data Center Through Automation

We're in a highly competitive environment and driving business agility has become job number one for every IT organization. Modernizing the data center is fundamental to powering agility and innovation, and automation is a major part of it.

Automation technologies help give organizations the flexibility to leverage any type of cloud, application, or infrastructure future. These innovative technologies are rapidly becoming table stakes for organizations seeking to achieve the speed and agility they need to meet the demand for faster access to IT resources—and enable a competitive edge.

Today, software business applications are emerging as the launch pad for new business models that are most responsive to escalating customer needs. Automation helps organizations accelerate IT service delivery by automating the configuration and provisioning of the networking, security, and infrastructure components that support application development.

In this paper, we will discuss how organizations can reduce complexity and eliminate repetitive, manual tasks by embracing an approach that fully automates the delivery and management of infrastructure and application components, including network and security services. We will also present some essential steps and best practices that organizations can use to prepare their infrastructures for an automation initiative.

Benefits of Automation

Automation enables organizations to eliminate time-consuming manual processes, siloed workflows, and risky, error-prone tasks, in exchange for more consistent delivery and management of IT resources. By automating infrastructure, IT teams can spend less time worrying about manual configuration and deployment processes, and free up time and resources to support innovation and business growth that they need to stay competitive.

Traditional hardware-based infrastructure is highly complex, siloed, and difficult to automate, requiring specialized, hands-on processes. Traditional infrastructures are also procedurally slow, requiring long lead times to provision IT and application resources. They rely on inefficient processes, which lead to unplanned downtime and wasted capacity. Security and compliance are also real concerns for traditional infrastructure environments, which offer minimal visibility into governance requirements, making compliance challenging.

Automation addresses the challenges of a traditional data center, using a software-defined approach that provides a single control plane across clouds, with access to native Cloud APIs that enable organizations to consume native services on any cloud. With integrated cloud management, IT can automate the delivery of composable infrastructure and app services, with self-service capabilities. Automation can also enable self-driving operations for continuous performance optimization, proactive capacity management, and intelligent troubleshooting and remediation, to enable IT to run production operations with more efficiency and less risk across private and hybrid clouds.

According to a recent paper by 451 Research, automation drove an 80% reduction in full time employee (FTE) expenses, 25% increase in productivity, and 30% reduction in downtime, as well as a 30% efficiency improvement in resolving incidents.¹

¹ "Can private cloud be cheaper than public cloud?," 451 Research, 2017

A Mindset for Automation

In the past, automation initiatives were very narrowly focused, and concentrated on improving very specific processes and highly targeted areas of the infrastructure. Today, in a multi-cloud environment, organizations must take a more holistic approach to implementing automation.

The first step involves corporate culture and communication. Organizations must work closely with business leaders and HR to refine job roles and create a company mindset where automation is not only welcome but encouraged.

Creating a Software-Oriented Team

The first steps toward automation focus on employee roles and skills. Existing IT operations teams should broaden their focus to take on more development responsibilities. In some ways, they will act as a Site Reliability Engineering (SRE) team, making a shift toward incorporating DevOps principles into their mission and responsibilities.

Instead of focusing on managing virtual machines (VMs) or hypervisors, organizations need an automation team that understands how to develop code. To best align networking functions with automation actions, IT may need employees with new skill sets, or might need to redeploy and reorganize its existing skill sets.

Gathering Automation Requirements

Before embarking on a full-scale automation initiative, IT should determine exactly what they are trying to accomplish, and which benefits they wish to achieve. They should consider the tasks and aspects of the infrastructure they wish to automate, and how best to serve their business stakeholders.

It starts with identifying repeatable processes and areas where automation can help pay off the “process debt.” For example, an organization might want to start with building the local automation capabilities first, such as automating the production process for a firewall change.

The team could consider the technical and process debt involved, and build code around it to automate this repetitive, manual process.

IT organizations should be sure their professionals are conversant with the varieties of coding languages required for automation, and are fully trained with the tools and platforms required to support their initiative.

With the right mindset, skills, and strategy in place, organizations will be ready to take the first steps toward automation.

Implementing Automation with the “Six Sevens” Best Practices

After careful evaluation of company culture and roles and responsibilities, organizations can embark on the automation journey. VMware has developed a set of best practices that are applicable to automating multi-cloud infrastructures, called the “Six Sevens”. This approach helps IT guide the underlying foundation for data orchestration and automation technology implementation.

The Six Sevens involves six process steps of proper data orchestration. These process steps support IT’s ability to understand and automate seven foundations of multi-cloud, discussed later in this paper. Let’s take a closer look at these Six Sevens best practices.

The Six: A Universal Multi-Cloud Data Orchestration Process

In a multi-cloud environment, most data orchestration scenarios can be boiled down to a set of basic requirements:

- A user will need to copy or move a specific partition of data, which may contain sensitive information, from a VMware cloud to a public or developer cloud, for a specific purpose.
- The output of the data processing is extracted and placed into production via containers.
- The original data is no longer needed and must be wiped on decommissioning.

These basic steps can apply to a variety of scenarios and environments. In some cases, the data may be in the form of a workload. This data movement can be broken down into six process components—a repeating pattern for making data available across clouds that is ideal for automation.

| THE SIX | |
|----------------------|---|
| 1. Data Description | Defining a data description is the first step in data orchestration. Data description focuses on understanding the nature of data and the effect of its loss, either through destruction or loss of security. Once organizations understand the nature of data, it can be partitioned, moved, used, or destroyed as needed for a variety of operations. |
| 2. Data Partitioning | Partitioning data involves setting up governance, security rules, or other guardrails for data that specify which extraction or aggregation can (or should) occur. For example, if an organization is seeking to extract part of a large data lake for training or other purposes, properly classifying the data will minimize risk and speed up the process. Data partitioning should occur as part of the process of preparing data for movement. The data description will play a key role in helping IT apply governance or security requirements. |
| 3. Data Placement | Once data has been sufficiently partitioned, it can be placed. This step includes selecting the right storage and security profiles—such as at-rest encryption and cloud endpoint—for the selected storage type. For example, if an organization is streaming data from a data lake cloud to a processing cloud, it may wish to set up local storage outside the data lake. Data placement should consider criteria like sovereignty, data gravity, and speed of data access. Again, a thorough data description process can help organizations better understand the nature of their data and how to place and process it. |

| | |
|----------------------------|---|
| 4. Connectivity and Access | All data movement requires some form of communications mechanism for its transfer. When thinking about connectivity and access, organizations should consider both network access and application programming interface (API) access. APIs can be used to help establish the right type of data storage and support for accessing it. |
| 5. Processing | Once data is placed, and storage, connectivity, and access are set up via networking, it's time to process the data. Machine learning (ML) is a key part of data processing. Organizations must define the nature of the code that will be executed, specifying criteria like workload type (VM or container), and the best tools for creating, deploying, configuring, and managing the data processing software. Cloud-native application models are not required for data processing, but they can help expedite the development and deployment of applications across multiple clouds. |
| 6. Cleanup | Once processing is complete, organizations will decommission cloud compute, storage, and network resources. Cloud operators often set up default decommissioning processes, such as disk wipes and network destruction, but default processes may not be enough to safeguard sensitive information. Organizations should consider regulatory and business compliance policies as they specify data cleanup processes. |

Now that we have discussed the six key processes involved in data orchestration, let's move forward and discuss what must be automated to make the process as efficient and secure as possible.

The Seven: Multi-Cloud Foundations

To automate the six processes of data orchestration that we have just discussed, organizations must closely consider seven key foundations to optimize operations across multiple clouds. These foundations are closely interconnected and include: data classification; governance; instrumentation; networking; platform-as-a-service (PaaS); security; and storage or public clouds, enabling not only business agility but also operational consistency and compliance.

| THE SEVEN | |
|------------------------|---|
| 1. Data Classification | <p>Classifying data creates a mechanism to describe data, and the effect of loss of any security objective. When automating multi-cloud data orchestration, organizations should consider issues like:</p> <ul style="list-style-type: none"> • What data is involved? • What are its qualities? • Is it growing at a high rate, creating potential future issues? • Are there legal or policy restrictions around its location? <p>Organizations should also consider how the description of the data influences requirements for storage, security, instrumentation, networking, PaaS choices, and other variables.</p> |

| | | |
|-----------|--------------------|---|
| THE SEVEN | 2. Governance | <p>Governance is the mechanism that lets IT form and monitor policies that control data access. Data access control also includes its movement, such as passing parameters to APIs. To determine how best to automate processes, organizations should consider governance questions such as:</p> <ul style="list-style-type: none"> • What government, industry, or business policies affect processing data? • Are there key auditing requirements? • What service levels are promised, and what are the impacts of missing them? • How are certificates, identities, authorization, and authentication handled on each cloud endpoint? |
| | 3. Instrumentation | <p>The ability to observe software execution is key to operations, and instrumentation enables it. Instrumentation provides mechanisms for projecting and sensing the operating characteristics of the processes (executable data) acting on other data or tracking its movement. Logging movement, network creation, access, and other processes is critical to understanding failures, as well as meeting compliance reporting requirements. Instrumentation solutions should encompass all processing metrics storage and analytics facilities, as well as logging services. When automating instrumentation, IT should ask whether their solution can:</p> <ul style="list-style-type: none"> • Observe the key service level indicators for all services. • Monitor for security breaches and ensure all data and processing are complying with policy. • Identify issues in the system and predict future failover or scaling needs. <p>To help safeguard the integrity of data and processes, logging services must also conform to proper security and corporate compliance standards.</p> |
| | 4. Networking | <p>Networking is the connectivity and access needed to move data across or within clouds. As cloud-friendly and cloud-native architectures at the software level have emerged, such as software-defined networking (SDN), and SD-WAN solutions, multi-cloud network setup has become more dynamically driven than traditional, manual networks. Nonetheless, networking will always involve a physical setup, increasingly by cloud providers. As organizations seek to automate their processes, they should consider networking questions like:</p> <ul style="list-style-type: none"> • Is sufficient network capacity available to run a workload in a given cloud or transfer data to a particular site, without causing unintended saturation? • Can the network ensure inflight encryption of data, based on governance and data classification? • Is the networking correctly implementing the needs informed by security requirements? • How do networking and dynamic provisioning processes for SDN or SD-WAN resources impact instrumentation, security, and governance? |

| | | |
|-----------|-------------|--|
| THE SEVEN | 5. PaaS | <p>Organizations can choose from a variety of PaaS platforms. To successfully automate their operations, organizations must consider their software delivery strategy and how it will support the processing stage discussed above. Some key questions to address when forming and implementing a PaaS strategy include:</p> <ul style="list-style-type: none"> • Are unique platforms required to accomplish processing goals and integrate specific applications? • Is the PaaS platform available in the clouds intended for use? • Can applications be tailored to a generalized PaaS to easily support multi-cloud deployments? Does the PaaS have built-in instrumentation, governance, security, and other services? |
| | 6. Security | <p>Security is all about assuring integrity of, and access to, data through micro-segmentation, encryption, and other technologies. It spans most every foundation and process across multi-cloud environments.</p> <p>When considering what mechanisms should be in place to automate the security layer of all data, organizations should ask themselves questions like:</p> <ul style="list-style-type: none"> • What level of security is needed for data in use, whether executable or content? • What are the consequences of compromising any key security objectives, such as availability, confidentiality, or integrity? • What methods are needed to detect and protect against compromise? • What are the policies for handling and preventing compromise? |
| | 7. Storage | <p>Storage encompasses not only basic criteria like disk availability, but also how data gets stored. For example, storing data in a Hadoop Data File System (HDFS) for use by ML algorithms also introduces data access and privilege issues. When evaluating storage, organizations should consider issues like:</p> <ul style="list-style-type: none"> • Are role-based access controls in place at the appropriate layer to assure proper limitations? • Does storage have the required speed or durability for the application? • Is the needed storage available in preferred clouds, or is a new provider necessary? • Does the storage need to be encrypted at the file system or block layer? • Does storage need to be wiped after purging of data? <p>Storage is inherently a fundamental issue that addresses multiple areas in an infrastructure. Organizations should consider its impact on security, governance, networking, and instrumentation—and how to automate those areas to properly use the storage without compromising data.</p> |

The VMware Advantage

It's clear that organizations in every industry understand the importance of business agility, and the potential of automation to enable this new level of capability. However, as automation becomes more widespread, it is also becoming apparent that the technology that enables it is in many ways the easiest part. An effective automation initiative is overwhelmingly based on people—including a company's culture, process, capabilities, and skill sets.

With a strategy and company approach in place for automation, organizations can dive deeper into the key process setups required for data orchestration, then consider the foundational elements of multi-cloud environments. With a best practices approach to considering these processes and foundation components, IT can minimize risk and maximize its return on automation.

VMware can help organizations use automation to achieve innovation and accelerate their time to market. Our comprehensive set of solutions enable enterprises to manage and run consistent infrastructure and operations across data centers and public clouds that deliver their most critical applications. VMware's solution is built on a software-defined HCI architecture of natively integrated compute, network, and storage virtualization technologies, with self-driving operations and management. This innovative approach delivers a consistent operational model on any cloud.

Pullin Innovative Technology Your Way

PROPEL YOUR BUSINESS FORWARD WITH VMWARE

Learn more about IT Automation >

Take a test drive in a Hands-on Lab today >

Join Us Online:



For more information contact: Pullin Technologies

info@pullintechnologies.com

678-592-2871

pullintechnologies.com

vmware®



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Copyright © 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: DCMA-0479_VM_Practical-Guide-to_Reducing-Complexity-with-Automation-Technology_WP