WHITE PAPER – **NOVEMBER 2018**

A MULTI-CLOUD PATTERN: THE SIX SEVENS

vmware®

Table of Contents

1. Introduction

Every week, VMware® professional services and sales organizations see more customers use multiple clouds as they transform into digitally driven organizations. While some have well-formed multi-cloud strategies and related implementations, gaps often exist between best practices and the reality "on the ground" as they progress along the digital transformation journey.

This article explores what issues underlie multi-cloud use best practice. In particular, it introduces a pattern for multi-cloud usage referred to as the six sevens. The six sevens involves six process steps of proper data orchestration to understand and automate the seven foundations of multi-cloud, discussed further below. Before detailing the six sevens, let's discuss the reason they exist in the first place.

1.1 Moving Data Across More Than One Cloud

There are many reasons organizations use multiple cloud endpoints. For example, pharmaceutical research companies may have to restrict access to drug discovery data assets on a tightly controlled cloud. However, analytical drug discovery engines may be on another developer cloud—a cloud supporting integrated platforms for the development, delivery, management, and maintenance of software applications and services.

Also consider businesses that operate so well that they can determine and shift to the least expensive cloud. That means not only determining the lowest real cost, but also that the data must be accessible or transferred cross-cloud. This involves moving executable data—such as VMs or containers—across clouds, cross-connecting clouds such as via a programmatically automated SD-WAN like VMware NSX® SD-WAN by Velocloud®, or possibly moving data underlying the code (source code and build jobs) to create the executable image on and for the selected cloud.

1.2 Each Cloud Endpoint Has a Purpose in the Value Chain

The examples above involve the use of public clouds. However, isolated yet elastically consumable units of compute, network, and storage also open avenues in identifying clouds that are not public. For example, consider a highly scaled Internet of Things (IoT) implementation—a worldwide network of uniquely addressable interconnected objects. This notion gave rise to fog[1] and edge computing[2], which distribute compute, networking, and storage closer to the end user (between the cloud and the things).

In whatever form, the clouds discussed have specific purposes. They carry out operations necessary to create added value from the information—data—flowing between. To amplify the notion of value-add in this case, consider that end-to-end IoT implementations normally follow a standard pattern of **ingest**, **analyze**, and **engage** (for example, engage with consumers by providing new information gleaned from the original data).

[1] M. Chiang and T. Zhang, "Fog and IoT: An Overview of Research Opportunities," IEEE Internet of Things Journal, vol. 3, no. 6, pp. 854–864, 2016.

[2] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge Computing: Vision and Challenges," IEEE Internet of Things Journal, vol. 3, no. 5, pp. 637–646, 2016.

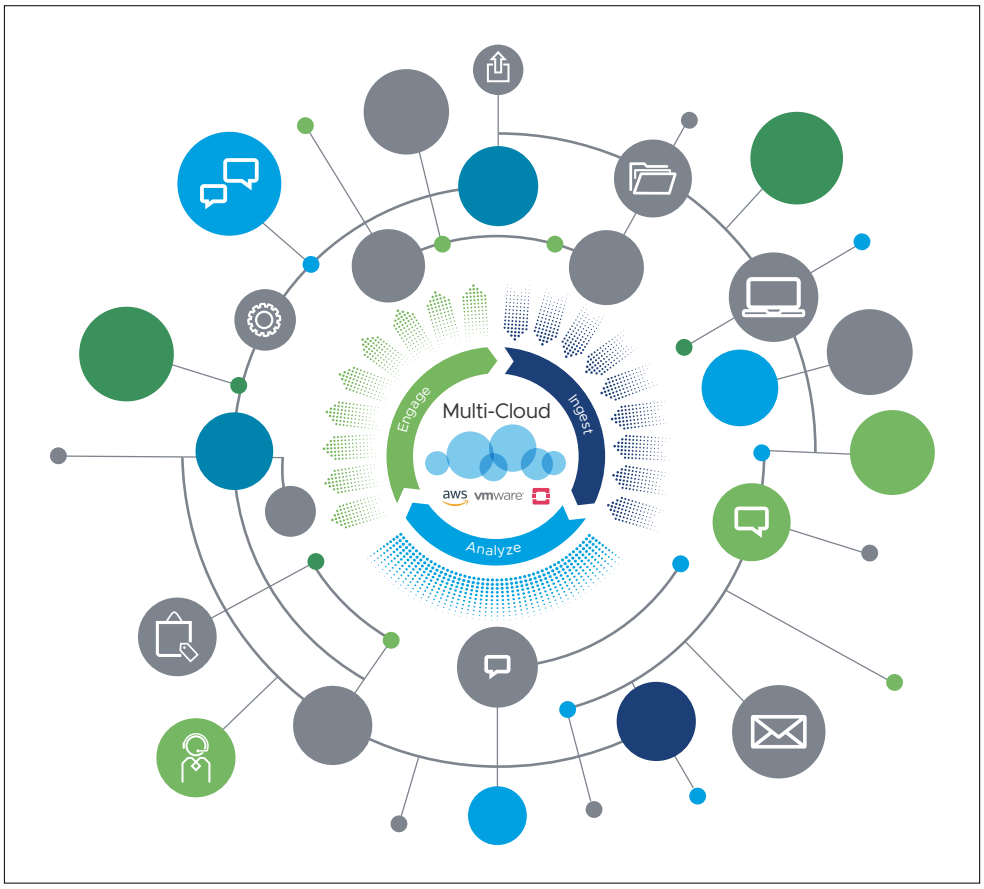

The center of the diagram above shows the general IoT pattern of ingest, analyze, and engage as a never-ending circle of activity. It's common to see edge clouds for ingest, aggregating sensor data from the external world. Fog or public clouds are used to analyze the ingested data and create deep neural networks (DNN) and other analytics[3] models. A combination of edge and public clouds are used to re-engage with the external world by executing the resultant DNN graphs as containers, orchestrated by an enterprise-grade Kubernetes implementation[4].

While not a simple operation by any means, it represents a classic case for creating a model for continuous customer engagement—a fundamental component of VMware's IT value model. For an example, see this demonstrative system that shows how customers can benefit from the ingest, analyze, and engage patterns as described.

The point is, whether IoT, AI-based drug discovery, workload placement strategies, or a plethora of other usage patterns, multi-cloud use is likely a need in the future of most organizations. Shifts are taking place in software development, where machine learning (ML) is becoming more prevalent, pushing single-cloud usage to extremes not seen in the past. Organizations must understand how they will successfully and safely utilize the broader multi-cloud ecosystem.

[3] Blog.dellemc.com, "Using a World Wide Herd (WWH) to Advance Disease Discovery and Treatment," December 2016.

[4] VMware.com, "Deploy Enterprise-Grade Kubernetes with VMware Pivotal Container Service (PKS)," December 2017.

2. Multi-Cloud Foundations: Distilling The Six Sevens Pattern

Because IoT implementations often operate across multiple clouds, IoT represents an easy way to peer into the six sevens pattern. Whether the edge and fog clouds are public or private is a matter of choice, but the fact that multiple clouds are in use requires care and consideration. The processes and data are always at risk of tampering and hackers. Data is at stake, as is the relationship with the customer who may be guided by the ML results produced, all of which is dependent on maintaining security objectives around the data flowing across various clouds.

2.1 The Six: A Universal, Multi-Cloud Data Orchestration Process

Data orchestration can be boiled down to relatively simple user story statements. For example:

"I am a [user] and need to copy or move [a partition of data 'X'], which contains corporate-sensitive and customer personal information, from my VMware cloud to a public cloud so we can execute DNN training. After training, I want to extract the resultant graph and place execute it in production via containers. Thereafter, I no longer need the training data, but need to be sure it is wiped on decommissioning."

By changing a few words here and there—for example "public cloud" to "developer cloud"—virtually all data movement user stories can be represented. Consider also, as discussed above, that moving workloads is akin to moving **executable data**, thus can be represented by the user story.

This relatively simple data movement story can be broken into six process components that form a repeating pattern per data that needs to be available across clouds.

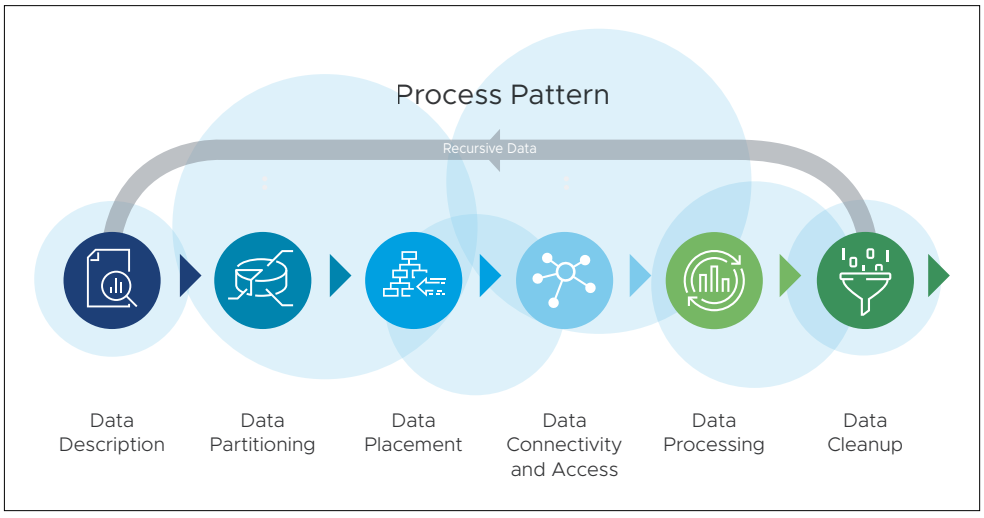

Note that as a repeatable process, data orchestration—much like (executable) container orchestration—should be an automated process. While automation concepts such as DevOps and site reliability engineering are beyond the scope of this article, such consideration should be given while reading the rest of this paper.

2.1.1 Data Description

Generally stated, in order to govern the use, exchange, transport, secure, or any other action on or about data, one must know the nature of data and its loss. The term "loss" as used in this context contemplates more than just mere destruction (an availability problem). Loss includes the sacrifice of any of three security objectives: availability, confidentiality, or integrity.[5] The loss of any security objective has consequences to both internal and external organizations.[6] Once the nature of data is understood, it can be partitioned, moved, used, and destroyed as needed for various operations as suggested by the user story.

For reference, the three security objectives noted are credited to the National Institute of Science and Technology (NIST), a non-regulatory, U.S. federal agency. They are relevant to other geographies[7], so they're a good starting point in any classification effort.

2.1.2 Data Partitioning

Data partitioning involves setting guard rails (such as governance and security rules) around which data access can or should occur. For example, in the user story provided, the goal is to copy part of a big data lake and use it for DNN training. While one can generally access data through various programmatic mechanisms, properly classified data will significantly expedite and safeguard the copy process.

Data partitioning should allow for data handling without unwanted or harmful side effects. The ML developer in the user story may seek to copy data for use locally on her laptop to independently work on DNN training models. Relevant data (in ML, the features and related known predictions) should be reasonably obtainable without the risk of unintended side effects. If the training data contains personally identifiable information (PII), the data movement will likely cause a loss of confidentiality, and thus breach of one of the security objectives. Proper partitioning can and should occur in the process of preparing data for movement.

2.1.3 Data Placement

Once sufficiently partitioned, placement of data becomes possible. Data placement involves, among other issues, selecting the right storage, security profiles, and cloud endpoint (location) for the selected storage type.

For example, streaming the data from the data lake cloud to the cloud analyzing it, storage may be required on the originating cloud to locally place and store the data for streaming outside the context of a data lake. If batch processing is intended, likely the data will be stored within the processing cloud, therefore requiring remote placement onto storage.

Such criteria raise issues of data sovereignty, gravity, and speed of data access. The nature of resolving these issues involves information from the data description process. Further, the placement is informed by the nature of the data and the processing it may undergo.

---

[5] See Gutierrez, Carlos M. & Turner, James M., "Guide for Mapping Types of Information and Information Systems to Security Categories," NIST Special Publication 800-60 Volume I Revision 1, National Institute of Standards and Technology, Gaithersburg, MD 20899, August 2008.

[6] id.

[7] iapp.org, "How NIST security controls might help you get ready for the GDPR," March 2017.

2.1.4 Data Connectivity and Access

Placing data creates the need for connectivity. By definition, data movement requires some form of communication mechanism for its transfer. Communications produce the need for—among other things—network access and application programming interface access (such as AI). Programming interface access can be called upon to create the intended storage and networks for accessing that storage.

2.1.5 Data Processing

Once data gets placed—implying that storage, connectivity, and access is setup, therefore networking exists—processing the data becomes possible. Processing in the user story involves ML. At a more mundane level, ML involves the same issues as any software: someone writes code that needs to be executed. The nature of that execution, such as processing, requires attention to issues like containment (VM or container) and the selection of tooling for the creation, deployment, and configuration and subsequent management of that software.

This gives rise to considerations of Platform as a Service (PaaS) solutions, such as Pivotal Application Service and VMware PKS. Cloud-native application models are not required, though they tend to expedite the development and deployment of applications across multiple clouds.

2.1.6 Data Cleanup

Once processing is complete, the decommissioning of cloud compute, storage, and network resources becomes necessary. Cloud operators often setup default decommissioning processes such as disk wipes and network destruction.[8] Regardless, if any snapshots of relevant disks exist, they may contain sensitive information. Therefore, disk wipes may not be enough. This represents a case where a recursion of the six processes applies to evaluating the data movement from disk to snapshot. Furthermore, cloud providers' disk wipe procedures do not guarantee the wipe process conforms to regulatory requirements or compliance policies that apply.[9]

Laying out the six processes involved in data orchestration leads to the discovery of what must be automated in order to implement the process as efficiently and safely as possible.

2.2 The Seven: Multi-Cloud Foundations

Investigation into customer scenarios similar to the general user story above has led to the conclusion that there are seven unavoidable concerns that require examination in order to automate the six processes of data orchestration. They are deeply interconnected, so it bears more fruit to look at them holistically as opposed to individually. The seven foundations of multi-cloud that must be automated in order to optimally operate across multiple clouds include data classification, governance, observability, networking, PaaS, security, and storage.

[8] Amazon Web Services, "AWS Security Best Practices," August 2016.

[9] id. ("When you have regulatory or business reasons to require further controls for securely decommissioning data, you can implement data encryption at rest using customer managed keys, which are not stored in the cloud.")

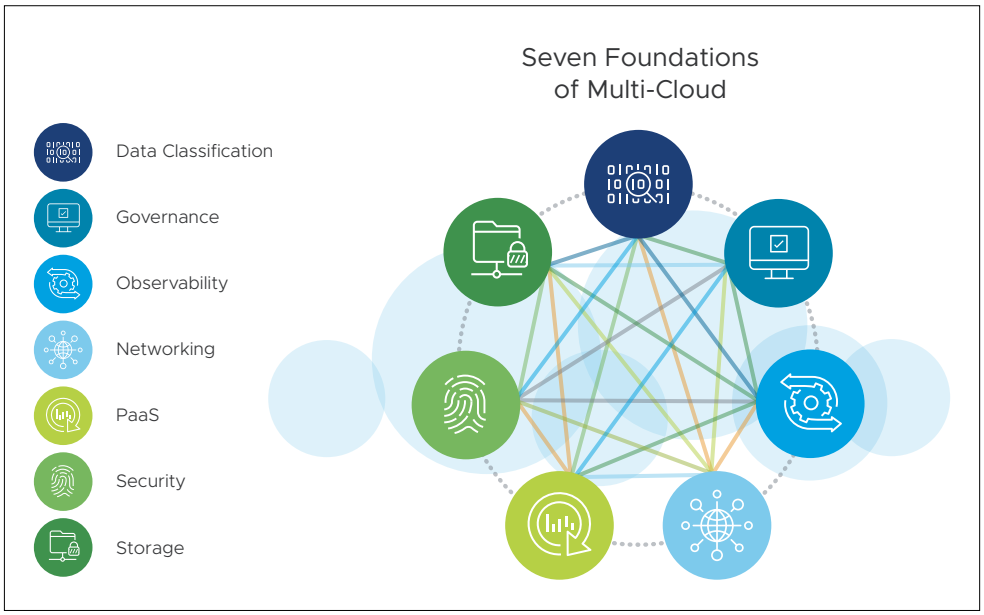

Note that the order is not significant. The foundations are highly intertwined in that each informs the others. In fact, by visiting any of the foundations listed, one will eventually visit all other foundations in order to complete an automated resolution of the six processes discussed above.

2.2.1 Data Classification

When automating multi-cloud data orchestration, a number of questions are answered by classifying the data. What data is involved? What are its qualities? Is it growing at a high rate, and therefore a future data gravity problem? Are there legal or policy restrictions around allowed placement of the data? What are the operational implications of a compromise of the data and can they be automated? Must operations on the data be time limited? In general, how does the description of the data influence requirements for storage, security, governance, observability, networking, or the choice of PaaS?

By no stretch of the imagination is this an exhaustive list of issues in order to classify data to safely automate data movement. The list is much longer, and the topic of following articles.

2.2.2 Governance

Governance is the mechanism by which organizations compose and subsequently monitor policies that control data access. Note that accessing data includes its movement, which includes accessing and passing parameters to APIs. Therefore, access encompasses a software program (i.e., data processing) or merely the movement of data to a screen so a human operator could view it. Interestingly, the reality is the latter case is itself a data placement issue – can the data be placed on the screen for those having immediate access to view? The answer to that question involves understanding the classification of the data and allowed placement policies.

All such movements involve resolution of many questions. What government, industry, or organizational 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 in and across each cloud endpoint? How do the choices within governance influence decisions concerning security, networking, and observability?

This non-exhaustive list of questions involves classifying the data itself, but also informs the other foundations, particularly networking, security, and storage. Regardless, a visit to all of the seven foundations may raise governance questions. Answers relevant to automated resolution are needed to inform the mechanisms.

2.2.3 Observability

In any scenario of software execution, observability is a critical component of operations. Observability involves providing mechanisms for projecting and sensing the operating characteristics of the processes (executable data) acting on other data. Remember that acting on data may merely be its movement from one cloud to another prior to ML training. Monitoring the effectiveness of that move—as well as keeping historical logs of the movement, network creation, access, and the like—is critical to not only understanding failures, but also proving the properness of the activity to regulatory auditors.

Observability solutions must support the incorporation of all processing metrics storage and analytics facilities like Wavefront™ by VMware, and logging services such as VMware vRealize® Log Insight™. By providing observability facilities, these solutions provide answers to questions such as whether you watching the key service level indicators for all services. Are you monitoring for security breaches? Are data and processing complying with policy? Are you able to identify issues in the system and predict future failover or scaling needs?

Observability itself should be part of every foundation involved. Logging, for example, must ultimately conform to proper security standards. It is not uncommon to find that, by accident, private keys or passwords are passed in clear text to logging systems, which themselves require storage. Is that storage encrypting data at rest, so as to minimize the possibly of information leakage? Such considerations help to understand that resolving observability is a fundamental issue involving and informing every one of the foundations.

2.2.4 Networking

Networking involves creating the connectivity and access required to move data across or within clouds. This field is certainly not new, but the nature of network creation has changed over the years. With the advent of cloud-friendly and cloud-native architectures at the software level, SDN (such as VMware NSX Data Center) and SD-WAN (such as VMware NSX SD-WAN by Velocloud), multi-cloud network setup has become more dynamically driven software operations than a hardware-level, statically defined networking.

With that noted, networking will always involve physical setup to support the dynamically driven SDN, but that is becoming the domain of cloud providers. Therefore this article does not address that other than to consider, somewhat implicitly, that physical access to networking hardware also requires care and consideration of data classification. For example, leakage of network appliance passwords or physical access to (smart) cards and racks clearly represents potentially significant loss of multiple security objectives.

Networking itself—dynamically driven or, if necessary, statically driven—requires consideration of whether sufficient network capacity exists to run a workload in a given cloud, or transfer data to a particular site without causing unintended saturation. Can the network ensure in-flight encryption of data, based on requirements revealed by visiting governance and data classification? Is the networking correctly implementing the needs informed by security requirements? What bearing does networking and processes that dynamically provision SDN or SD-WAN resources have on observability, security, and governance?

Again, this is not an exhaustive list of questions, but gives an indication of considerations raised in order to assure no loss of any data security objectives for SDN, SD-WAN, and any cloud network API.

2.2.5 PaaS

The processing stage of the six processes involved in multi-cloud data orchestration requires a reasonably formed software delivery strategy. PaaS provides many basic needs of software delivery. There are many PaaS platform forms available today and each has its pros and cons.

Some key questions to address when forming and implementing a PaaS strategy involve considering whether unique platforms must be used to accomplish the processing goals. For example, does the PaaS provide for facilities that expedite, help secure and implement ML, IoT data ingestion or SaaS (e.g., SalesForce.com) application integration? Is the PaaS platform available intended for use in the cloud? Can the PaaS platform and application be tailored to easily support multi-cloud deployments? Does the PaaS have built-in observability, governance, security, and other fundamental services?

In general, consider how the PaaS influences observability, governance, security, and storage.

2.2.6 Security

Security involves assuring the integrity of data. Security is a broad topic that includes micro-segmentation, encryption, and the like. Security affects nearly every foundation. For example, should a process be allowed to execute on a given VM, such as being controlled by VMware AppDefense™? One can consider this executable block similarly to a port block on micro-segmentation implementations. This is just one of a myriad of security considerations that one must take into account when transporting data—executable or otherwise—across clouds.

One issue to consider is the level of security needed for data in use, whether executable or content. What are the consequences of compromise, where compromise means the loss of any of the three security objectives (i.e., availability, confidentiality, and integrity)? What methods are necessary to detect and protect against compromise? What are the policies for handling and preventing compromise? In brief, what mechanisms must be in place to automate the security layer of all of the data within the multi-cloud environment? How does security inform storage, observability, data classification, networking, PaaS, and governance?

2.2.7 Storage

Storage—for example disk availability—is often viewed as a rather low-level issue. However, storage as contemplated in the seven foundations involves not only raw storage, but also how data is stored. For example, storing data in HDFS for use by ML algorithms forces the consideration of data access across the entirety of the distributed file system. What users are allowed to access which files on the HDFS? Are role-based access controls in place at that layer to assure proper limitations?

Questions of interest are ones that consider the need for storage. Does it have the required speed or durability for the application? Is the needed storage available in preferred clouds, or is a new provider (i.e., data placement process) necessary? Does the storage need to be encrypted at the file system or the block layer? Does it need to be wiped or scrubbed to a certain degree after the purge of data?

Storage is inherently a fundamental issue that addresses multiple areas: hardware, file systems (software), encryption and therefore key rotations and so on. Considerations include how storage informs security, governance, networking, and observability, and therefore how to automate those areas in order to properly utilize the storage so as to prevent any loss of the data security objectives.

3. Conclusions

VMware is seeing organizations frequently try to leverage multiple cloud providers. Many of these organizations are struggling with not only the complexity of cloud usage, but the processes and mechanisms to automate multi-cloud use. Because these organizations are undergoing their digital transformation, a multi-cloud strategy will likely dominate IT operations for—or at the very least in—the foreseeable future. VMware has laid out the six sevens pattern and related consulting services to help with this extremely important part of the digital transformation.

VMware offers assessments and gap analysis services that give a clear view of the current state as it relates to the six sevens, as well as roadmaps and designs for resolving any gaps for the future. The teams involved help identify best practices for multi-cloud operations including setting policy, training resources, and guiding staffing requirements. The anticipated outcome of these services is to help our customers transition existing and new workloads into a multi-cloud strategy that is well thought out and efficient to operate.

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