Technical White Paper–2021

Sovereign Cloud

TECHNICAL WHITE PAPER (VERSION 1)

JULIAN DA SILVA

Contents

Figures

Introduction

VMware Sovereign Cloud Framework is a set of guiding principles and best practices for delivering cloud services that can meet security, compliance, and data sovereignty requirements for a specific jurisdiction in which that cloud operates. This white paper discusses the technical design and operational considerations when designing for a VMware Sovereign Cloud using a subset of content from the framework. Although this white paper is primarily technical, a collaborative approach to design is recommended to facilitate design decision making that ideally involves both technical architects and security specialists or data compliance officers.

The information provided in this white paper is for illustrative purposes only and is not intended to constitute legal advice. VMware encourages customers and partners to obtain appropriate advice from qualified advisors for the interpretation of all relevant data privacy laws, industry regulations, compliance requirements, and any other applicable requirements that apply in jurisdictions in which they operate.

Purpose and Overview

This whitepaper discusses an approach to designing a VMware Sovereign Cloud based on established VMware reference architectures, describing strategies for overlaying technologies and design patterns to develop compliant and secure compute platforms that host and process data while maintaining sovereignty. Furthermore, a functional Sovereign Cloud platform must be supported by appropriate operational processes and procedures that maintain and enforce the secure and compliant status of sovereign data, workloads, and the platform as a whole. The final design will depend on the compliance and legal environment in which the platform will be deployed and the service offerings and target customers the platform is intended to serve. This white paper does not stipulate which services must be offered or how they must be configured to be compliant with specific security frameworks. It does propose suggested services, features and characteristics that are needed to develop compelling Sovereign Cloud service offerings that represent value for customers and help to address security and compliance requirements.

Target Audience

The information provided in this white paper describes technical design considerations and options that architects and engineers can consider together with specialists who are professionally responsible for interpreting IT security compliance regulation and applicable law within their jurisdiction. In some cases, this may be the same person or team such as a solution architect with specific expertise in secure Government IT environments. However, in many cases these roles will be separated and so this discussion is tailored to address design options in the context of risk and mitigations.

Foundational Knowledge

The core concepts that apply to a VMware Sovereign Cloud platform are described in this section as well as foundational technology concepts related to VMware Software Defined Data center (SDDC) and the VMware Cloud Provider Platform (CPP).

VMware Sovereign Cloud

The core proposition of a VMware Sovereign Cloud is to provide enforcement of data and workload residency, extend data sovereignty protections beyond the immediate platform where possible and to enable secure, audited connectivity and data transaction between resident, sovereign and non-sovereign data classifications.

A VMware Sovereign Cloud intends to support the cloud provider with the following customer profiles:

- Hosted Government IT systems, data, and applications
- Hosted Military IT systems, data, and applications
- Private sector organisations and academia working directly with Government (linked IT)
- Private sector organisations and academia storing and processing data of national interest or sensitivity
- Customers across all industry sectors seeking platforms that are data privacy and security centric

In all cases, a VMware Sovereign Cloud provides the tools needed to enforce data privacy and security restrictions, auditability, assign classifications to data assets, offer comprehensive value-add services for industry and Government, and a trusted platform for hosting both traditional and modern workloads flexibly and reliably.

Data Guardians

Maintaining the sovereignty of customer data requires careful consideration as to how to continuously manage data classification and lifecycle, this applies to data as it is processed as well as when it is transferred and stored at-rest. This potentially expands the role and responsibilities of the provider as well as the customer data owner and implies a greater degree of operational cooperation.

To enact and oversee these additional responsibilities, providers might consider assigning the role of data guardian to one or more support individuals. A data guardian acts as an advocate for the customer whose data resides in the provider environment, providing continuous oversight of the provider’s conformity to agreed data security and management principles. A data guardian might also offer an advisory service to customers to recommend data protection and handling strategies appropriate to the data classifications in question that make best use of the provider’s offerings depending on the need. Services might include, encrypted and immutable storage, database tokenization, sovereignty assured DR, and data protection services in which all service elements reside within the jurisdiction.

Such a role would be reflected technically in a Sovereign Cloud platform by way of role-based access control as well as in operational processes to support the platform. A data guardian would require customer-relevant audit data and even access to platform controls for the purpose of system quarantine, alert notifications, and other operations depending on how mature the provider’s sovereign services are. In a similar manner, providers who offer advanced security services might be expected to directly intervene in customer environments in reaction to a cyber-attack based on pre-agreed plans and authorizations, a data guardian could act to prevent unintended data breaches.

Security Domains

When designing for a secure, compliant cloud platform, security domains become the first consideration before other more traditional considerations such as performance, scalability, etc. Security domains are the starting point for the design and all other design decisions flow from this point.

In broad terms, a security domain is a conceptual grouping of systems, network connections and supporting infrastructure that fall within a common security boundary. One way to look at security domains is that they represent an area of common trust in IT systems, ensuring that all participants are subject to the same security and trust criteria. For the purposes of a VMware Sovereign Cloud, the definition is expanded to people and operational processes.

Security domains typically represent a common authentication and authorization boundary (e.g., an LDAP realm), such that being granted access to one system in a security domain leads to the explicit or implicit granting of access to other systems in that domain. The degree of access will vary based on access control mechanisms, firewall restrictions, or other security components that are implemented. In many cases, being granted broad access to many systems at once in a security domain is desirable depending on the function of the authenticated individual or application. When this is not desired or if the scope of access must be restricted to a subset of systems within the domain, then access control mechanisms must enforce restricted scopes, or alternatively those systems are placed in a separate security domain. Similarly, a security domain can represent a network connectivity area with a common security posture with protections located at the domain boundary, security domains can optionally be further subdivided into smaller connectivity areas using subnetting and micro-segmentation.

Architecturally, every IT system represents one or more security domains, these are reflected in their design both physically and logically, the latter in particular. Common examples of security domains in IT environments found globally include Management, DMZ, Production and Test & Development.

Many logical security domains are represented directly in software and are enforced in a product's architecture whereas others are established dynamically using scoping and grouping functionality configured by the customer or provider to define boundaries to which controls, policies and rules for access and interaction are applied.

An example of a logical security domain represented in software would be an Org entity in VMware Cloud Director which is the product's representation of a tenant in the cloud platform. The creation of an Org results in the automatic creation of an inventory folder within vCenter (management plane of the virtualization platform) in which the tenant's workloads will be placed, acting as an administrative grouping. The Org contains a locally scoped user account database within Cloud Director and is also routinely configured with an external directory for centralized or federated authentication of the customer to the Org. This action results in the Org becoming an administrative and security boundary that aligns a customer with their cloud resources with dedicated authentication against a designated authentication source.

An example of a dynamic logical security domain would be a security group in NSX-T for the purpose of network security. Virtual machines, networks, and other entities can be placed in an NSX-T security group so that a common set of firewall rules can be applied to all members. Membership of security groups does not have to align to other logical groupings in the platform and can span compute clusters, sites, and other constructs if desired. This is useful when defining a common connectivity area between different parties, a new security domain with its own access controls in the form of firewall rules can be established independently of other scopes in the platform.

They key difference between this example and the Org example earlier is that the scoping of Orgs is useful only for the creation of a tenant construct, whereas a security group's membership is flexible which makes them useful for implementing various strategies depending on the need. There are many dynamic scoping constructs available in a VMware platform that are useful in representing logical security domains and enforcing common security policies.

Overall, all member entities of a security domain generally fall under a common security posture which is maintained through policy control. Policy-driven security is essential for managing a large number of entities in an automated way, for reducing administrative burden for administrators and change approvers, and for reducing the possibility of human error.

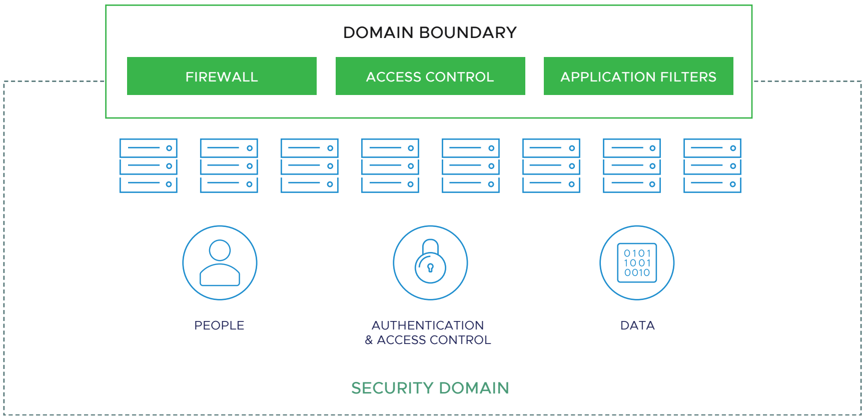

Figure 1 - Conceptual Security Domain and Boundary

All VMware Sovereign Clouds must include two prescribed security domains: A Resident domain and a Sovereign domain. These domains encompass both management and workload domains in vSphere as well as all supporting infrastructure and management elements.

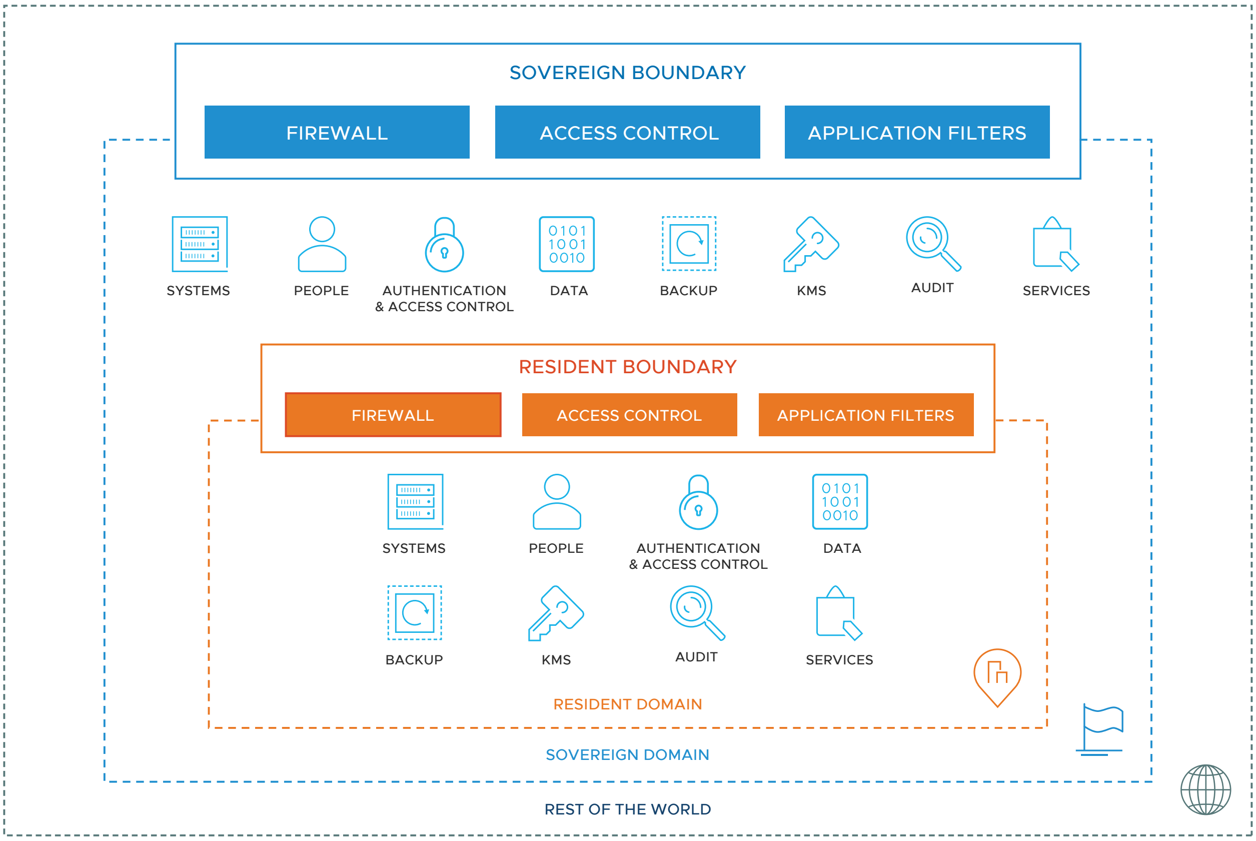

Figure 2 - Initial Sovereign Cloud Conceptual Security Domains

The Resident domain represents an area with the highest degree of trust and confidence where all member elements reside within the geographical boundary of the jurisdiction. No data is stored or processed outside of the jurisdiction and no network traffic ever traverses hardware that is not assured to be within the local geography. No element of the domain can be deemed to be under the control or influence of foreign or non-sovereign entities. The resident domain only accepts traffic from systems in its parent Sovereign domain and other qualified Resident domains within the same country or pan-national entity, it has no route to and from the public internet or other provider networks. External networks that are deemed highly secure and sovereign, such as Government and Military private networks can be presented in the Resident domain.

The Sovereign domain represents an area of reduced trust and higher risk mitigation in which network connectivity between member systems and non-sovereign systems is possible on an exceptional basis. This domain serves use cases in which the customer application cannot be entirely isolated and is required to interact with external and on-premises systems. Storage and processing within the Sovereign domain would primarily still be resident in the local geography, however, the Sovereign domain does allow for the possibility of replication or data transfer outside of the jurisdiction but with enhanced protections such as encryption at-rest and in-flight. This principle might be compared to a nation state's overseas Embassy, in which all the laws and customs of the 'home' state prevail within the Embassy walls but while being physically present in a foreign jurisdiction. In the Sovereign domain, systems and data are exposed to a theoretically higher risk but are subject to superior protection to mitigate that risk.

An example usage pattern for an application might be a Government Public health website which is internet-facing and that requires access to a database of confidential personal records on the back-end. In the Sovereign Cloud architecture, the database would be hosted in the Resident domain and would be accessible only to the website's application servers. The application servers would be hosted in the Sovereign domain as this is the only domain in which a workload can be reached from the outside while at the same time be able to connect to Resident systems. In this example, it could even be the case that the website front-end is hosted in the provider's commercial cloud platform which in this architecture is represented as the 'Rest of the World'. If this approach is used then the sovereign status of the application servers is extended to the web front-end hosts which, although exist outside of the Sovereign domain, are assured to have an equivalent level of operational governance.

It should be noted that Resident and Sovereign domains are not reflected as specific application constructs in VMware products, they are instead reflected by Cloud Director provider VDCs and other logical objects visible to customer and providers in the platform.

Cross-Domain Interaction

As all member entities of a security domain share a broadly similar security profile and area of connectivity, the rules and policies that govern their interaction are equally similar. Of course, there will be exceptions for specific workloads and traffic flows within a security domain, but permissible connectivity and authentication patterns for member entities will often be applied in common. Where security domains will interact with one another, these interactions will take place at a perimeter or boundary in which security controls are placed in order to protect the security of domain members by restricting access and connectivity against more stringent rules. For example, network communication traffic is necessary between systems in a management security domain and systems in other domains. As the management systems are sensitive and have the potential to impact operations across multiple systems, inbound network connectivity is usually limited by firewall rules to clients used only by administrative staff and to ports and protocols used by specific management agents that might be deployed on systems in other domains. These restrictions serve to maintain the integrity of systems in the management domain and to reduce the risks and impact associated with attacks, malware, and ransomware. Security control and mitigation functionality such as firewalls are placed at the domain boundary. The conventional approach is to place these on or adjacent to network gateways as this is at the point of network traffic ingress and egress for the domain. Mitigation is an important consideration as sometimes it is not possible to entirely eliminate security risks for systems that must be reachable. Allowing any kind of connectivity exposes a potential attack vector. This requires that strategies and techniques are put in place to greatly reduce the risk of an attack and limit the potential for harm should an attack take place. Monitoring activity at the security domain boundary is key to successfully identifying and thwarting attempted and successful security breaches. This results in prompt and timely reactions such as blocking connectivity or even automatically rebuilding systems that are suspected of being breached.

The required level of control and mitigation functionality depends on the level of trust for systems within that domain. The DMZ environment is a security domain in which connections are accepted from remote systems that cannot be trusted. A DMZ is typically composed of security hardened systems with very limited access to trusted network environments or even adjacent systems in the same environment and will not have direct access to the same internal resources such as a directory service or file storage.

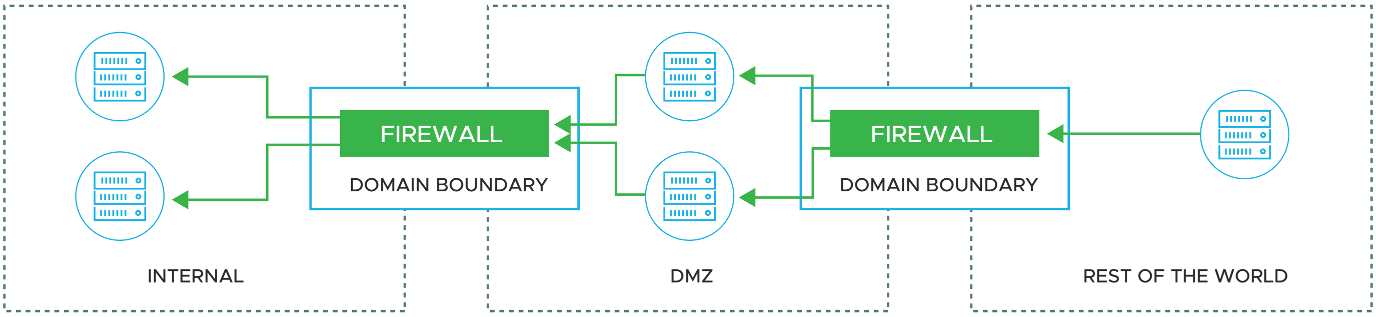

Figure 3 - Cross-domain Network Connectivity Flows

In a VMware Sovereign Cloud, the concept of the DMZ is augmented into a Sovereign security domain in which the scope of untrusted remote systems is not limited to those on the public internet but to any system (including other provider hosting platforms and management systems) outside of the immediate platform. This is discussed in more detail later in this white paper.

As security domains are logical constructs it is entirely possible for one domain to be nested within another and for different domains to overlap.

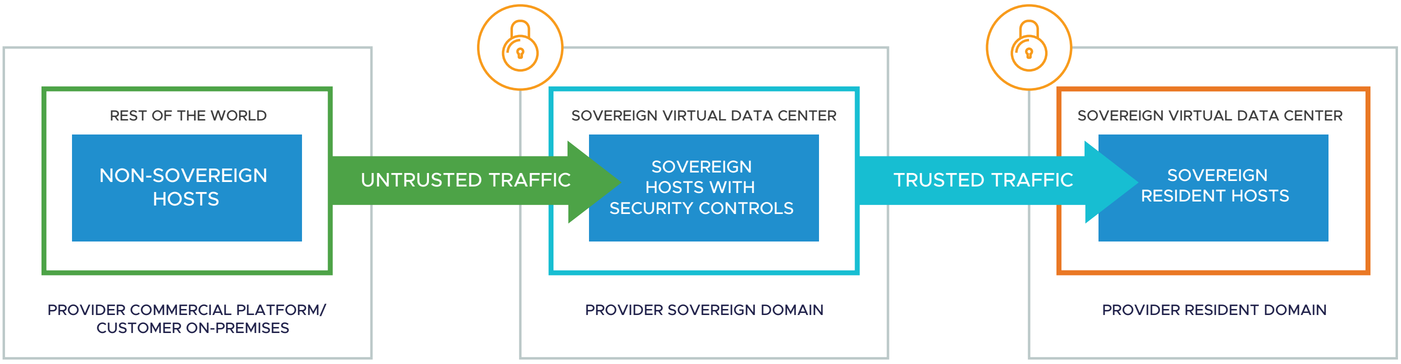
REST OF THE WORLD
NON-SOVEREIGN HOSTS
PROVIDER COMMERCIAL PLATFORM/ CUSTOMER ON-PREMISES
UNTRUSTED TRAFFIC
SOVEREIGN VIRTUAL DATA CENTER
SOVEREIGN HOSTS WITH SECURITY CONTROLS
PROVIDER SOVEREIGN DOMAIN
TRUSTED TRAFFIC
SOVEREIGN VIRTUAL DATA CENTER
SOVEREIGN RESIDENT HOSTS
PROVIDER RESIDENT DOMAIN

Figure 4 - Sovereign Cloud Domain Trust Transition

The network and data flow patterns between the outside world and the Sovereign and Resident domains transitions from un-trusted high risk connectivity scenarios to highly trusted low risk connectivity scenarios.

The following rules are specified for interaction between the Resident and Sovereign domains in a VMware Sovereign Cloud where the Sovereign and Resident domains are to be implemented:

- Systems in the Resident domain can never connect to non-Sovereign endpoints, this includes anything on the public internet and anything outside of the Resident or Sovereign domains in the platform, including other platforms hosted by the provider
- Systems in the Sovereign domain can only interact with systems in the outside world that themselves have elevated protections including encrypted network connectivity and data encryption at-rest. The exception to this is if additional risk mitigations such as application proxies and IPS are implemented in the Sovereign domain boundary
- Resident domain networks must not be routinely stretched or bridged into the Sovereign domain, this should be done only in exceptional circumstances and be protected by additional network security measures if implemented
- Implement zero-trust, micro-segmentation in the Sovereign domain for all customers
- Make micro-segmentation available in the Resident domain, zero-trust does not necessarily have to be enforced
- Always encrypt network traffic for sensitive data classifications before it leaves the domain boundary, ideally all traffic. The simplest way to enforce this is to configure the provider firewall of the domain boundary to only allow VPN traffic to ingress and egress the domain. By ensuring that the top-of-rack switches and domain boundary firewall are physically secured and locked within the rack where sensitive data is stored this means that no data can enter or leave the rack in an unencrypted state

- A Resident domain can span multiple geographic locations if all connecting elements can be verified and assured to be under complete control of sovereign entities and reside within the same geographical jurisdiction. This includes the physical sites that the domain will span, the links that provide connectivity between the sites and the land (and waters) that the links traverse. Government and Military site links are typically redundant with diverse routes and protected by Military-grade cryptography. This level of protection may be sufficient to satisfy the needs of customers in the jurisdiction of the platform, if not then a dual-site Sovereign Cloud platform is achieved by splitting each site into its own Resident domain and regarding the inter-site links as being part of the Sovereign domain. Architecturally, this would invoke the operating rules of Resident/Sovereign domain network and data interaction. Workload migrations and replications would be considered as a transition between domains with all the security and risk mitigations that would apply. This scenario should ideally be avoided due to the potential for adverse impact on operations but may represent a cheaper solution overall if cost is a priority
- Favor the use of provider-controlled private IP schemes and discourage the use of customer public and private IP address schemes, such as from secure private Government networks. Controlling the IP scheme for customer workloads provides opportunity for less disruptive migration if workloads are migrated between the Resident and Sovereign domains
- Resident domain north-south traffic must not connect to the provider's core network. The Resident provider-layer gateway must only be uplinked to gateways in the Sovereign domain or to Sovereign private networks such as Government and Military networks
- Implement change control and approval processes for firewall changes at each domain boundary and any operations that involve the migration, duplication, import and export of workloads in a domain. As well as involving conventional stakeholders in the approval process such as change managers and application owners, also include those who's remit specifically involves oversight of data protection and sovereignty, such as a data guardian or security officer
- Avoid the use of 1-to-many and many-to-1 firewall rules for north-south traffic flows that involve IP ranges or object groupings. Instead, specify 1:1 mappings so that each firewall rule isolates specific hosts
- If desirable, implement a NAT and Load-balancer approach for IP presentation and avoid routing from the Resident domain towards the Sovereign domain. This allows for original IP addresses to be obscured which is useful in environments where metadata leakage is a concern
- Implement a black hole default route on the Resident domain's provider-layer gateways so that all IP routes must be deliberately specified in order to enable up-stream connectivity. Normally, a network gateway will direct all traffic it does not know how to route based on its own routing tables to a default gateway address. By implementing a black hole default gateway at the domain perimeter, traffic is dropped if the required route is not explicitly configured by the provider

Controls & Mitigations

A VMware Software Defined Data center (SDDC), represents a platform for hosting both modern and traditional applications with comprehensive security controls and features that are essential for a secure cloud hosting platform. Additional VMware solutions offer security-focused capabilities that further augment the platform's ability to secure both itself, the applications, and data it hosts. These can be complimented by additional 3rd party vendor and open-source solutions that may be integrated in various ways. Examples of additional guidance that are available include the VMware **Validated Design Security and Compliance Configuration for ISO 27001** and the **Security and Compliance Configuration for VMware Cloud Foundation**. Always check with vendors for the latest available guidance.

Authentication & Authorization

Every transaction and operation in a secure computing environment must be attributed to a known and trusted entity whose activities are authorized and traceable to ensure accountability and integrity of data and systems. Authorization is assigned by an authority either directly (approval for a requested activity), or in most cases, through preauthorization by way of an individual account with assigned roles and permissions that persist for a predetermined period. Proving that a person, system, or application is authorized to access resources starts by demonstrating with an extremely high degree of confidence that the request for access originates from the trusted entity the request claims to be from. Strategies for authenticating entities range from the use of account/user and password combinations to multi-factor authentication that includes the use of passwords and other methods of authentication such as certificates and security tokens simultaneously, 2-factor (2FA) most commonly.

Ensuring that these principles have the desired effect, a rigorous approach towards configuring the technology and designing the processes that govern how people and systems interface is needed.

The following recommendations for implementing trustworthy authentication and authorization in a Sovereign Cloud environment are made for the components that fall under the provider's control which includes but is not limited to VMware components:

- Never allow anonymous access on any management interfaces, this ensures maximum traceability
- Do not expose customer-facing management interfaces, such as self-service portals, to the public internet. Access to these interfaces should be protected by a VPN. Also consider the use of IP geo-filtering to prevent attempted connections from outside of the country to block authentication events from foreign entities
- Never use shared accounts. Every person, application or system must be assigned an account exclusively. Avoid using common service accounts for specific functions such as backup or monitoring agents. If functionality exists in these applications to provide alternative credentials for each target system, then make use of it
- Implement a complex password policy with requirements for minimum length, compulsory inclusion of both alphanumeric and special characters, prevention of historic password re-use and expiration. Ensure that all systems that include their own local account databases have these policies set to the greatest extent that the system supports
- Use centralized authentication against a secure, trusted directory service to centralize access management and authentication audits
- Implement signed certificates on all interfaces that support them to provide a high degree of trust that a connection made to another system in the platform is the intended target endpoint. Sign and issue certificates from a certificate authority that resides either within the Sovereign Cloud platform itself or within another sovereign compute environment that the platform is connected to that is under the complete control of the provider and/or customer. Issue certificates that expire in less than 2 years or as short as is operationally practical.
- Never allow clients and management tools used to manage the platforms that are capable of supporting the certificate authority's chain of trust to connect to an endpoint with an invalid or self-signed certificate. In these circumstances the operational procedure should be to reject the connection and to remediate the issue (if the endpoint is ultimately found to be trustworthy)
- Always set an expiry on accounts, even for permanent staff, customer accounts and system accounts. For the latter, extra care is needed to ensure that service account passwords are well managed to ensure that there is no interruption to service. Make use of reminders and alerts to ensure that action to change passwords is taken with sufficient time to allow for change control and support staff coverage in a planned change window
- Related to the previous recommendation, routinely review accounts and validate that they are still needed. Delete unused accounts that no longer serve a purpose. Re-confirm the status of administrative staff to ensure that they maintain their vetted status (where applicable) and that their circumstances have not changed in any way that might require their vetted status to be withdrawn or re-qualified
- Do not re-use passwords for active accounts from one account directory in another account directory or application. In secure hosting environments it is common to have to authenticate separately to systems with differing security status, there is a temptation to simply re-use the same password to make life easier, this must be avoided. Implement these operational procedures and consider introducing an offset in password policies between environments such as differing minimum password lengths and password ages
- Provider directory services must be entirely under the provider's ownership and control including all other identity sources that the directory service instance is federated with. In an authentication context, federation involves trusting specific authentication providers (such as another directory service) in a different security domain to authenticate entities within the domain. All components in the directory must be hosted entirely within the local geographical jurisdiction, this extends to backups and remote replications
- Avoid authenticating against local account databases that exist at a system or application level. These typically exist to enable first-time configuration of an application before being commissioned or for troubleshooting purposes when centralized authentication is not possible. Ensure that local authentication events are centrally audited

- Always generate complex passwords using random generation, consider using a dedicated enterprise-grade credential manager for this purpose but ensure that it is hosted within the platform itself
- Require individuals to set a new complex password under the same password strength policy the first time they log in. This ensures that no person other than the rightful account holder knows the password while the account is active and useable
- Assign extremely long and complex passwords to system accounts that are impossible to remember and highly impractical to take note of. As these accounts are not for use by people then the opportunity should be taken to make brute-force attacks completely impractical
- Implement a policy of password cycling for service accounts to generate new passwords each time they are used by administrators
- Consider implementing a policy of leaving administrative accounts disabled by default and only enabling them when a ticket has been raised that relates to a system or application for which the administrative account is needed
- For highly secure Government or Military environments, consider implementing a split password approach whereby service account passwords are not revealed in their entirety to one individual but are instead split between two individuals who then must enter their part of the password in the appropriate sequence. This approach prevents authentication by lone actors, and which requires a witness to be present during authentication and subsequent activity
- Always implement 2-factor authentication for individuals accessing the management plane of the platform. Ideally use security tokens in combination with passwords. Consider using client certificate-based authentication to prove the authenticity of the client
- Implement 2-factor authentication for services and applications where possible
- Assign accounts to cloud provider and personnel only after carrying out the required verification of identity of the individual, vetting of the individual based on their role and confirmation of their legal citizenship. Ideally, vetting and identity verification will be conducted by an independent party, this may in fact be a legal requirement in some jurisdictions
- Use usernames/account names that do not include any characters, abbreviations or phrases that offer any clue as to their purpose, who they are assigned to or what function they support. Use random sequences that are not difficult to remember but reveal no information as to their purpose

VMware Cloud Director and vRealize Automation both support Microsoft Active Directory and OpenLDAP-based directory sources, please refer to product documentation for specific guidance and version support. In addition, both products support authentication using SSO which uses the SAML 2.0 protocol. This protocol is used by many identity providers, enabling the use of multifactor authentication and other enhanced security capabilities offered by 3rd party vendors. Both products offer multi-tenancy capabilities and allow providers to use different authentication schemes and directory sources for different tenants if needed.

The overarching principle for authenticating entities in a Sovereign Cloud and authorization of entities to interact with it is to guarantee that no element of the authentication or access control process is hosted or takes place outside of the sovereign jurisdiction. If any element, such as an LDAP directory server or a federated directory entity exists outside of the jurisdiction then this breaches the sovereign status of the platform. Note that this has implications for multi-national organizations that implement global authentication platforms. At the Sovereign Cloud platform level, global authentication platforms cannot be used. They can, however, be used by customers for the authentication of their hosted workloads in very specific circumstances.

Access Control

Access control relates to the permissions, rights, and privileges that an authenticated entity holds in a system. The following access control models are the most prominent today and are summarized here:

- **RBAC –** The most common model is role-based access control (RBAC) in which the terms of an entity's ability to interact with systems, applications and data is prescribed by an assigned list of granular permissions allowing specific actions to be performed that are appropriate to the role of the authenticated entity. As the granular permissions available in many applications can be extensive and finely detailed, they are usually bundled into groupings referred to as roles and are assigned by way of administrative group membership. This is the model that will most likely be employed by providers for implementing access control to administrators and customers, it is implemented in all products in VCF and the Cloud Provider Platform
- **ABAC –** Attribute-based access control (ABAC) is based on the principle of granting access based on specific characteristics of the authenticated entity defined by stated attributes. This approach to access control introduces additional flexibility in how access is granted but requires technology by a 3rd party that also supports the use of SAML 2.0 SSO

VMware products used by cloud providers implement roles which support separation of duties for both provider administrators and end-customer consumers when interacting with the platform. For a more detailed discussion on what roles are available by default and how additional roles can be implemented please refer to the official documentation for the relevant VMware product.

The following recommendations apply to the operation of a Sovereign Cloud:

- Implement the principle of assigning least privilege, granting accounts only the minimum required rights and permissions to perform a job or service function
- Provider engineers should have separate accounts for privileged and non-privileged operations, using the privileged account when elevated rights are needed to perform an activity and the non-privileged account for all other times such as updating tickets in a ticketing system or sending unclassified email. In a highly secure computing environment, such as a Military or intelligence environment then this principle may be extended to using different accounts for different security classifications, missions, or end-customers
- If practical, limit the number of customers that an engineer can interact with, assigning them to specific customer accounts for specified period. Consider rotating personnel between customer accounts on a random basis and without giving the engineer advanced knowledge of which customers they will work with
- Consider implementing automated assignment and withdrawal of privileged access rights to provider engineers over customer resources as and when they are needed such as when the customer raises a support ticket. Provider engineers would have no routine access whatsoever to customer resources until it became necessary. This strategy requires that the ticketing system itself is highly secure and fully audited and cannot be administered directly by support engineers
- Design and implement a comprehensive range of roles for the platform that support separation of duties for provider administrators and customers. The number of roles will depend on the operational process design of the provider, appropriate to the customers and security classifications served by the provider
- Related to the previous recommendation, create roles dedicated to audit and observation of the environment. Various products in the Cloud Provider Platform offer pre-defined roles that effectively offer read-only visibility of parts of the platform for the purpose of security audit and observability
- As a matter of routine process, review an entity's assigned roles and permissions to ensure that those held are still appropriate to the role. Remove any permissions or access to systems that are no longer needed

Multi-Tenancy

Providing cloud services to multiple tenants using a shared compute platform is the most common operating model as it offers efficiencies in cost, operations and floorspace that can be passed on to customers through economies of scale. It is achieved predominantly at the logical level whereby different tenants are abstracted away from the underlying infrastructure and from one-another in software. Another model that is frequently offered alongside the shared compute platform model is to have dedicated compute hosts assigned to specific customers with many supporting infrastructure elements such as the management plane and elements of the physical network infrastructure remaining shared. Using a combination of these approaches is commonplace in commercial cloud environments where the priorities of the end-customer are usually cost and performance. High availability and fault tolerance capabilities have long been offered from VMware-based shared platforms and do not necessarily require dedicated hardware. However, for nationally sensitive Government, Industry, and Defense customers, physical separation of workloads from other customers can be a security and compliance requirement. This also extends to the management plane and supporting infrastructure. The need for additional hardware represents a cost premium and an additional administrative overhead, however, opportunities exist for service providers to make better use of logical tenant segregation, partly thanks to advanced security functionality but also because of having a platform with an overall superior security posture – 'a rising tide lifts all boats.

Adding multi-tenancy capabilities to a VMware platform involves introducing a management software layer, delivering Cloud Management Platform (CMP) or Service Management (SM) capabilities for customers to interact with and multi-tenancy functionality to segregate them from each other and from provider-managed elements.

VMware provides two CMP solutions to provide out-of-the-box cloud management functionality to a VMware SDDC. Products by other vendors may also be placed in-front of these CMPs to introduce an additional service management and governance layer (check which versions of VMware software they support).

Caution is advised when selecting a solution to provide secure multi-tenancy as some products only offer a 'soft' form of multi-tenancy that is optimally suited to environments in which tenants are part of the same legal entity, such as different departments in an organization. The degree of separation needed in a Sovereign Cloud must be suitable for 'hard' multi-tenancy for environments in which tenants have no direct legal relationship.

- **VMware Cloud Director -** The most straight-forward approach is to deploy Cloud Provider Platform which incorporates VMware Cloud Director (vCD). Cloud Director creates the logical abstraction from the underlying infrastructure and segregates tenants into Organizations which act as security and administrative boundaries that isolate each tenant's visibility and control. Customers are allocated Virtual Data centers (vDCs) that represent a unit of compute, storage and network resources that are securely dedicated to that customer, hiding elements of the infrastructure the customer is not entitled to. Integration into the underlying vSphere platform ensures that compute and memory resources allocated to customers are enforced so that one customer can never consume more than they are entitled to and cannot impact the platform's performance for other customers.
- **VMware vRealize Automation –** An alternative CMP product with significant functional overlap with Cloud Director is vRealize Automation (vRA). vRA has been historically employed by enterprises to assist them in transforming their IT operations from a traditional IT delivery model to a private cloud delivery model. Using vRA, customers can blueprint infrastructure and platform services and provision instances of these services onto both VMware and non-VMware hypervisor platforms.
- **3rd Party Tools –** Some cloud and infrastructure management products are available on the market that provide service management and orchestration capabilities as well as multi-tenancy abstraction. These tools can be overlaid directly over vSphere and in some cases can integrate directly with VMware Cloud Director and vRealize Automation.

Data Protection Strategies

The protection of data includes preventing unauthorized access, modification, and duplication as well as to prevent both intentional and unintentional loss or corruption. The VMware Cloud Provider Platform and the wider VMware Tanzu portfolio include a range of capabilities for the ongoing protection of data when hosted on a VMware-based platform.

Hosted customer data can be protected using various methods depending on the level of protection required and this in-turn depends on how the data is stored and accessed.

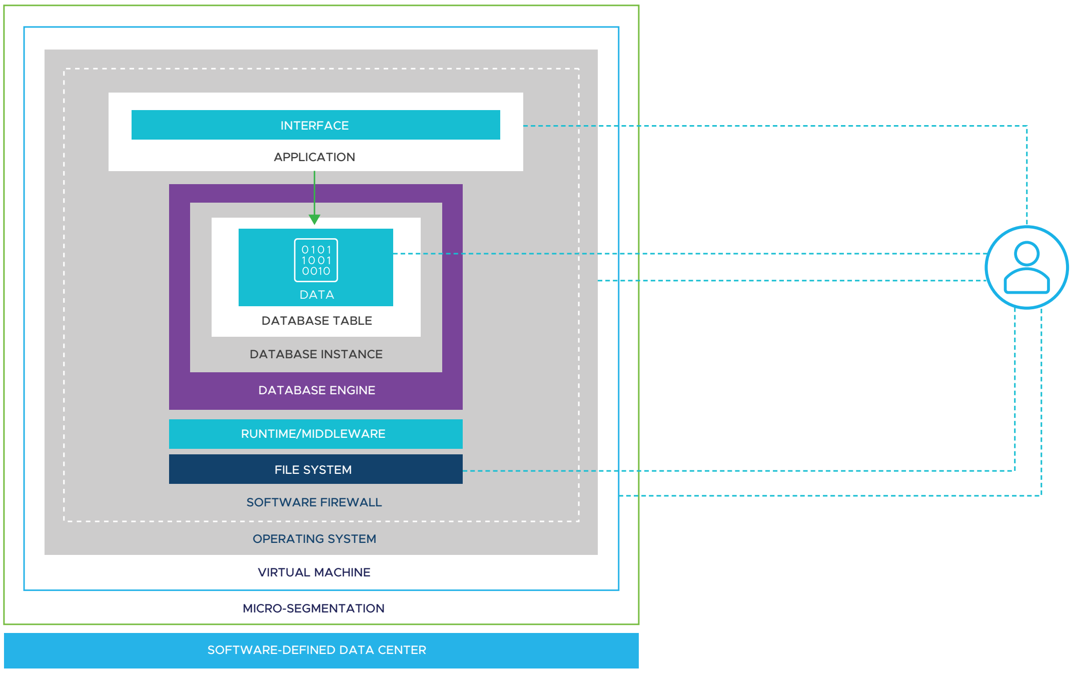

Figure 5 - Various Interfaces for the Access and Control of Data

Another important consideration is how the protection of data is impacted by any data lifecycle strategy that applies to the data. Data duplicates, including backups and replicas count as discreet data assets. Data assets themselves can be version controlled, assigned access control, and sharing rights. These assets may be subject to specific intellectual property protections or to legal rights of the data subjects or owners. These characteristics do not change when applying data protection strategies. Making a backup of a data asset creates additional versions of the data that should be tracked in the same way as the primary data asset and any metadata or access controls that apply to that data should be preserved such that they can be restored along with the primary data itself. An important element of data sovereignty is that when the data owner chooses to delete the data, that this is acted upon. So tracking data set copies and versions throughout the lifecycle is necessary.

To assist both the provider and the customer in protecting data and managing it in a compliant way, it is recommended that the data's ownership, categorization, and protective status is tracked using metadata. This metadata information should be available to all parties who are responsible for handling or controlling the data, which would ideally also be available to platform automation. There are various tools and strategies available to assist providers and customers to achieve this depending on how the data is stored and processed.

A catch-all strategy for protecting data involves securing the platform so that access to data is restricted. Encrypting all storage protects hosted data in its entirety at-rest, however, taking this broad approach may have undesirable performance and cost implications for certain data classifications which makes this a critical design decision. Encryption of data at-rest can be applied at multiple levels in the application stack. As well as encrypting the storage, it is necessary to encrypt all network traffic carrying sensitive data where it ingresses and egresses the platform by way of a VPN, this requires a suitable VPN client on the other end of the connection.

In reality, different data protection strategies can be employed depending on the sensitivity of the data being protected, the legal and compliance requirements of the customer, jurisdiction, and the impact of a data breach. The objective of the Sovereign Cloud provider is to offer the required level of protections to customers while also seeking ways to manage cost and making the platform as flexible as possible to consume.

Network Security

Securing network traffic is necessary both within the data center environment in which the platform is hosted as well as for end-to-end connectivity scenarios. A breach of network security at any stage risks invalidating the sovereign status of the data the platform hosts. Consider the following recommendations when designing a Sovereign Cloud:

- **Physical –** All network cables, switches and other network hardware appliances connected to the Sovereign Cloud should be physically secured and made inaccessible to any person who is not specifically authorized to operate the Sovereign Cloud. This physical protection applies within the data center which is already presumed to be secure and where other non-sovereign services are hosted. This also applies to customer and provider owned hardware on which network links are terminating for connectivity into the Sovereign Cloud. All hardware related to the Sovereign Cloud must be isolated in dedicated data halls, cabinets or cages in their entirety, and these must be locked. This approach ensures traceability of actions within the Sovereign Cloud to known and authorized entities. The only network elements that should be exposed to other areas of the data center are uplink connections and devices that link the Sovereign domain of the platform to the outside world

- **Firewalls –** A VMware Sovereign Cloud will already include edge and distributed firewall capabilities within the platform, the provider will likely include additional firewall appliances, possibly at the insistence of certain customers. As a minimum, the Sovereign Cloud will include an edge firewall in the Sovereign domain which represents the primary ingress/egress point for the entire Sovereign Cloud, an edge firewall in the Resident domain to control traffic flow between the Resident and Sovereign domains and an edge firewall for each tenant virtual data center. Some customers may introduce additional firewall layers using virtual appliances or physical firewall appliances hosted in the provider data center

- **Micro-Segmentation –** NSX-T provides an additional layer of network security with its distributed firewall capability which enables secure micro-segmentation of the platform. Micro-segmentation provides firewalling capabilities that are applied at the interface level. Network packets are inspected and filtered as they enter or leave an interface on a virtual machine, providing maximum protection and the ability to enforce zero-trust network environments. Providers and customers can use micro-segmentation to create restrictive scopes of network communication appropriate to their application and apply firewall rules directly to those groups. Micro-segmentation with zero-trust enforcement is a requirement of the Sovereign domain, it ensures that no host can communicate with another host without being explicitly allowed by a firewall rule. Zero-trust is a requirement of the Sovereign domain as this is the only domain that allows direct connectivity to external systems and is therefore an attack surface. Zero-trust mitigates the risk of lateral attack in the unlikely event of a host being compromised from outside. Zero-trust is not a requirement of the Resident domain, however, this feature should be available to customers should they decide to implement it

- **Network Isolation –** There are multiple layers of network isolation recommended in a VMware Sovereign Cloud, customers are free to introduce additional isolation if desired. The Resident domain of a Sovereign Cloud is isolated from the outside world; no inbound connections are possible thanks to two layers of edge firewall and a lack of routing to enable access from the outside world. For a Sovereign Cloud, the outside world refers to any system or network that is not part of the platform itself where total security equivalence cannot be assured. This includes the public internet but can also include the customer's own on-premises networks and the service provider's own internal systems and networks if these have not been secured physically and logically in equivalent ways. One useful advantage of this topology is that in the event of a suspected or confirmed attack, the link between the Resident domain which hosts sovereign data and the Sovereign domain where the attack lands can be severed – effectively raising the drawbridge between the data in the back-end and the application in the front-end. Overall application and data isolation is achieved by preventing unrelated applications from communicating with one-another in the Sovereign domain and access to data sources in the Resident domain being constrained only to the appropriate application

- **Encrypted Network Tunnels (VPNs) –** NSX-T provides both IPSEC VPIN and eVPN functionality for use by customers on a self-service or ticket basis and for providers to secure traffic flows between environments.

Encryption

VMware vSphere includes virtual machine storage encryption capabilities as part of the platform. NSX-T also offers network traffic encryption capabilities using IPSEC VPN functionality for both providers and end-customers to use. These capabilities can be used in conjunction with other capabilities offered by 3rd party vendors and the open-source community.

vSphere Virtual Machine Encryption

Encrypting data at-rest can be achieved at the virtual machine level in vSphere and is applied using storage policies. Encryption is applied to the virtual disk files attached to the virtual machine as well as the memory swap file and snapshots. Note that the virtual machine configuration file itself is not encrypted which includes MAC address metadata. vSphere Encryption relies on a compliant KMS server to generate and manage keys, vCenter requests the private key for a virtual machine when it powers on the virtual machine and so the KMS becomes a critical dependency, the key is not needed again until the virtual machine is again powered on.

vSphere Virtual Machine encryption is useful if the end-customer is prepared for the provider to manage the encryption keys on the KMS server on their behalf. Alternatively, if the customer procures a dedicated vCenter and vSphere cluster then the customer can provide and manage their own KMS server that the provider then configures the vCenter to use. This scenario allows the customer to hold the keys as well as the KMS infrastructure to gives them the ability to withdraw access to the KMS and therefore the ability to power on and access their virtual machines at any time.

For more information on How vSphere Virtual Machine Encryption works please review How vSphere Virtual Machine Encryption Protects Your Environment

Agent-based Disk Encryption

An alternative to vSphere Virtual Machine encryption is agent-based encryption using software from a 3rd party vendor. These products work by installing an agent on the guest that encrypts the file systems of the virtual machine and is managed by a management server. These solutions have no dependency on and do not interact with vSphere Virtual Machine encryption. Using agent-based encryption enables customers to own and operate their own disk encryption solution independently of the cloud provider. This is especially useful if the customer wishes to host some of their workloads outside of the Sovereign platform that will interact with the workloads that they have hosted in the provider's Sovereign Cloud. This would extend the sovereign status beyond the platform into multi-cloud and on-premises.

Storage Platform Encryption

Providers can use encrypting storage platforms to secure their data. This requires specialist hardware and additional licenses from the storage platform vendor. Encrypting at the storage layer encompasses array-based encryption and/or fabric-based encryption, with the latter securing data as it traverses the SAN fabric. When using array-based encryption, data is stored in encrypted form on the storage medium by the array. This has to be used in combination with fabric-based encryption to ensure that data is encrypted from the moment it leaves the VMware host to when it reaches the storage array. The use of fabric-based encryption may not be so important if the entire storage fabric is enclosed within the same rack or cage as the hosts and storage array, however, this may present a challenge at-scale.

Data In-Flight

The encryption of network traffic between security domains should be a matter of routine in a VMware Sovereign Cloud. By far the easiest way to achieve this is using NSX-T IPSEC VPN tunnels. These can be configured by customers on a self-service basis via Cloud Director. It is also recommended that traffic between a customer's Resident networks and Sovereign networks also be encrypted but this may represent too much of a performance burden for some customers.

Customer network traffic can also be encrypted at the application level using SSL. Most application servers include this capability and customers should be encouraged to make use of it. This should also involve using CA signed certificates and avoiding the use of self-signed certificates for anything facing the public internet. Providers can offer customers SSL offload functionality using NSX. The provider may choose to restrict inbound connections at the Sovereign domain boundary to ports traditionally associated with encrypted traffic such as 443 or may employ protocol filtering to block unencrypted traffic flows.

Application Proxy Filtering

The use of application proxies can improve application and data security by introducing additional layers of access control, auditing, and content filtering. 3rd party proprietary and open-source application proxies are varied and offer support for specific applications and protocols, some even provide traffic tunnelling functions. Application proxies can be an effective way to augment applications with new security, audit, and access control capabilities that the application itself does not natively support.

When placed between a client and a database or application server, an application proxy can inspect requests, responses, and act to modify or block transactions based on a policy. They also offer additional audit and monitoring capabilities. They can filter database requests based on data classification and other configurable factors. In the context of databases, a database proxy filter can inspect SQL queries, responses, and can allow, reject, or filter the content, such as to specific columns in the response to a query.

An example use case is a user who has permission to read all columns in a particular database. They can query all values using a client within the same country but if they make the same request from another country then the filter removes sensitive columns that require sovereign protection or blocks the request entirely. In this case, having permissions to access the data is not sufficient and the fact that the request originates from outside the nation is enough to prevent access.

If the application proxy has a public API, then this provides an opportunity for providers to use custom orchestration to create an integration with NSX-T to manipulate firewall and micro-segmentation policies based on access policies defined in the proxy.

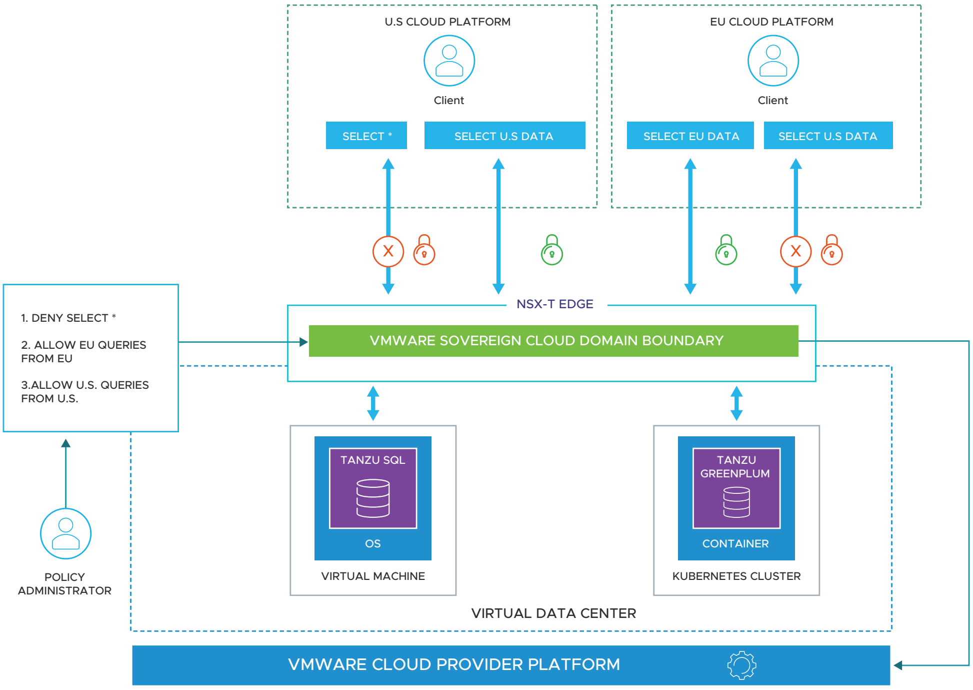

Figure 6 - Augmented and Enhanced Data Access Control at the Sovereign Domain Boundary

Data eco-systems consisting of multiple platforms and cloud providers, such as GAIA-X, will depend on specialist application proxy filters, known as connectors, to support sovereign data exchange and protect the sovereign boundary of the dataspace eco-system. These connectors securely encapsulate data transfers between authorized parties and grant access based on a usage policy. The function of the VMware Sovereign Cloud in this scenario is to extend the architectural principles of GAIA-X to the compute, network, and storage layer, and restrict data flow

An example of an application proxy filter is Envoy which is an open-source project. As well as including a range of features useful for securing connectivity and access to data, it is extensible and allows for new capabilities to be engineered.

Data Masking

Data masking is the process of obscuring the value of data so that it can only be read by users who are authorized to un-mask it. Characters in protected fields are substituted by cipher characters to hide the original content. It also offers a method to restrict visibility

This approach is similar to encryption but considerably less sophisticated in that it makes the data unreadable to those who do not possess a key or transformation method to reveal it. However, this is a simpler method that usually offloads the process of un-masking the data to the client application or user, it can also lack the strength of cryptography meaning that a determined attacker in possession of the masked data could eventually determine its value.

In some senses, this method of data protection is more complicated than outright encryption as it must be implemented at the client/application level.

Products such as VMware Tanzu Greenplum can be augmented to provide data masking capabilities, contact your VCPP Account Manager to find out more about VMware Tanzu Data Services.

Data Tokenization

The process of data tokenization involves substituting specific data values with unique tokens that reference the original value stored in another data source. This means that the primary data set can be made available to a client by the data controller without any private data values being present but is still useful as the private data is referenceable but unknowable to the client of the primary data set and retrievable by the data controller. An example would be to substitute a customer's name in a database with a unique customer ID in the primary data set and then create a secondary data set that contains both customer IDs and customer names. An authorized entity who is allowed to have visibility of customer names may use the secondary data set to look up the name of a customer using a customer ID, something that other entities are not entitled to do.

Tokenization offers opportunities for data owners and controllers to securely segregate data sets based on classification and desired protection levels and assign access to those data sets only to appropriate stakeholders.

Immutable Storage

Immutable storage technology ensures that data, once stored, cannot be modified, or deleted. This capability is useful in environments where trust in the integrity of data is of very high importance, typically for compliance and legal reasons. Once the data is stored it becomes a permanent record that cannot be altered by either the data owner or the provider as this is prevented by the technology itself. Note that the concept of immutable storage can apply to primary storage, backup solutions and removeable media. Customers may require immutable storage in any of these forms. Consider also that immutable storage technology does not in itself protect from loss in the event of a hardware failure and so redundancy may be necessary whereby the data is replicated to another instance of immutable storage, possibly off-site. When selecting an immutable storage solution, ensure that it is 'write once, ready many' compliant (WORM compliant), some WORM compliant technologies also allow for the configuration of pre-set time limits on immutability. The storage device or media itself must also be physically secured to prevent direct tampering or destruction.

Backup and Restore

In addition to offering virtual machine-level backup solutions, it is strongly recommended that application-aware backup solutions are also provided as they generally offer more granular backup and restore capabilities for applications and databases. This is useful for rolling back recent transactions while preserving data that was created after a backup was last made. Restore operations can also typically be executed non-disruptively as the application can remain online. The choice of backup solution(s) will to some extent depend on the applications and platform services that the provider offers to customers, examine the application portfolio being offered to customers and select a backup solution that offers the desired level of application support coverage.

DR & Replication

Disaster recovery is an essential capability of a cloud hosting platform being used by customers who require Sovereign Cloud services. The nature of the data and applications being hosted and their importance to the nation means that customers will want to recover quickly and reliably in the event of a disaster. For this reason, Sovereign Clouds should offer cross-site protection, replication, and fail-over services within the jurisdiction to give their customers the tools they need to achieve their business continuity goals and to preserve sovereign data integrity. VMware provides two products that provide DR protection for customer workloads: VMware Cloud Director Availability and VMware Site Recovery Manager (SRM).

A VMware Cloud Provider Platform can include integrated virtual machine DR protection using VMware Cloud Director Availability, presented directly to customers via a self-service UI. Customers can manage workload protection and failover or migration themselves, implementing tiered DRaaS SLAs and integration directly into VMware Cloud Director and the underlying VMware infrastructure.

Providers who opt not to use VMware Cloud Director or Cloud Director Availability can use SRM can to offer workload protection and recovery, integrating at the vSphere level.

The replication of virtual machine storage can be achieved using the integrated VMware vSphere Replication functionality that will work with both VMware products or the provider may opt to offload cross-site storage replication to a supported underlying storage system. It is essential that providers choose the appropriate storage platform when designing a DR solution and ensure that vendor licensing costs have been taken into consideration.

The design of a DR solution in a VMware Sovereign Cloud should consider the different security domains and data classifications that the platform is composed of. This may mean implementing multiple DR solution instances in a single platform to maintain isolation between different security classifications. The replication traffic between sites must be encrypted. This may be best achieved using physical cryptographic devices and the connectivity between the sites should be redundant, take diverse routes, and not depend on any entities that are controlled by foreign entities.

Separately to a DR-specific offering, make data replication functionality available to customers at the storage, database, or application level so that data of national importance can always exist in duplicate, if needed. Customers can be offered replicated storage volumes for the purpose of virtual machine and container persistent storage placement while services such as DBaaS can be configured in clustered or replicated multi-node configurations depending on the database technology involved.

Software Updates & Patches

An essential aspect of platform security is to maintain software and firmware that is current and up to date, allowing for testing before rollout to production. Cloud hosting providers depend on hardware and software vendors to notify and release patches to customers, VMware releases patches online for customers to download.

The following recommendations apply to providers building a Sovereign Cloud:

- Operate a model platform to minimal scale that includes all the functional components of the production environment
- It remains the responsibility of the provider to verify that any software patches will not have an adverse impact on the operations of their IT systems. Develop a rigorous and thorough test plan for evaluating patches and software updates that incorporate tests for inter-component integrations for both VMware and VMware solution components
- Install patches and updates gradually over time throughout a production environment as opposed to deploying them to all applicable systems at once. This ensures that any unforeseen adverse impact that was not identified in the provider’s testing have reduced overall impact
- Only download patches from official VMware repositories and verify their signatures against those we publish on our website, reject any patches and software packages that to do not come from VMware directly or who’s signatures are invalid or cannot be verified

  VMware’s Security Response Policy can be found **here**

Audit & Compliance

A key proposition of a VMware Sovereign Cloud is the ability to provide enhanced audit and compliance reporting, the output of which should be made available to customers on-demand and to official accreditors if this is a requirement in the jurisdiction. In the case of official accreditation, this may have to be periodically renewed and so it is important to continually gather supporting evidence. Security events such as suspected or attempted breaches should be reported to customers as well as the actions taken by the provider in response to those events. This not only provides openness to customers around the effectiveness of the provider's security measures and procedures but enables them to adjust their own security strategies in response to what could be repeated and persistent attacks. Ongoing audit helps to provide assurance to customers that their data remains protected in a compliant manner.

The VMware Cloud Provider Platform includes products that can assist in gathering audit information and provide continuous monitoring of the environment.

vRealize Operations (vROPS) is a monitoring and capacity analytics product that can be expanded to include compliance management packs that monitor metrics relevant to compliance in various industries. In addition, to compliance monitoring, ongoing capacity, and resource utilization monitoring, and alerting, using these capabilities helps to prevent outages and service degradation by warning in advance of eventual capacity exhaustion. vROPS is a requirement of a VMware Sovereign Cloud.

vRealize Log Insight (vRLI) is a centralized logging product that collects, filters, and helps analyze log output from VMware products and syslog sources, helping to bring all diagnostics and audit data together in a central repository. It is a requirement of VMware Sovereign Cloud to implement vRLI and to configure all VMware components for use with it. This is especially important for components such as Cloud Director and NSX-T as one handles customer-facing self-service activities and the other handles network connectivity and security.

Consider using VMware vRealize Network Insight (vRNI) for performing network traffic pattern analysis and audit on a one-off or continual basis. vRNI can identify traffic flows between both virtual and physical network nodes, build profiles of 'normal' activity, and identify unusual or concerning network traffic. In a secure network environment, routine monitoring and alerting of customer application network traffic events can represent a premium service that the provider can offer. This is useful not only for routine network security monitoring but also for pre-migration assessment of network dependencies considering a transition from a provider's commercial hosting platform to a new Sovereign Cloud offering.

Using output from these products as well as other tools from 3rd party vendors and the open-source community gives providers a rich and comprehensive source of detailed audit information that can be used to generate audit reports for specific customers and independent auditors as needed.

Trust vs Risk vs Cost

Most information security compliance frameworks take an approach that seeks to balance trust and risk against the cost of mitigations. IT operators design technical and operational mitigations and controls to manage identified risks as trust levels decrease, but by doing so the overall cost of the solution tends to increase. The level of risk that can be tolerated will range depending on the nature of both the data and the owner or controller of that data. Data controllers of a confidential customer account database may choose to tolerate a greater degree of risk than the Military might over IT systems of national importance.

When working with risk-based security frameworks, risks are identified and catalogued, and their potential impact is assessed. Protecting the sovereignty of data involves having a very low tolerance for risk and so various security frameworks will recommend highly prescriptive control processes and security functionality to be present and enabled.

Overall risk is ultimately relative to the total impact and liability of a breach. Damage is not just measured in financial terms but also in terms of its impact to the national interest, this is just as relevant to private contractors who provide services to Government agencies as it is to Government and Military entities themselves. Risk should be assessed openly by all parties involved. In cases where the sovereign boundary is being extended beyond the platform, the customer should take some responsibility by managing risk for the assets under their own control (such as on-premises infrastructure, VDI terminals, support personnel, etc.)

A breach of the sovereignty hosted on a cloud platform could have the following implications:

- Financial and reputational damage to the provider
- Financial and reputational damage to the customer
- Damage to interests of the nation and its allies
- Damage to national reputation
- Interference with sensitive systems and data
- Theft of intellectual property, including technological, scientific, or Military innovation

Identify available mitigations and include them in the architecture design and operational process design, attempt to determine the financial costs associated with those mitigations. Consider both improbable and probable risks. Then determine which risks will be mitigated, which risks will be rejected (i.e., risk will not be taken), and which risks will be accepted without mitigation. Ensure that a risk register is maintained and made available to customers. It may be useful to have separate risk registers against different security classifications as the risk profile of each will differ.

Data Residency

Data residency refers to the physical location in which data is stored and processed, by extension this includes the systems on which that data resides. Controls within a Sovereign Cloud platform over workload placement, data storage, and processing ensure that the resident status of the data is always preserved. In some limited cases, these security controls can enable placement of data on other sovereign platforms where residency is assured, such as to another Sovereign Cloud instance in a Military data center that is reachable over an encrypted link.

Data Sovereignty

Data sovereignty refers to data being subject to the laws of the jurisdiction it is collected in. This generally implies that the data is generated within the geography of the jurisdiction and that it will remain there, but this does not necessarily have to be the case if the sovereign status of the data is maintained. Many organizations around the world store and process confidential data in other countries in which there is a high degree of confidence that the data will be stored and processed in a manner compliant with the laws of its originating jurisdiction. In these cases, the legal protections afforded to that data in its home jurisdiction are extended with legal accountability for breaches falling with the data processor or controller in the home jurisdiction. This is due to the mechanisms of public international law that provide normalization across borders in many countries and include data protection and privacy within their scope. This brings about situations in which even companies that are entirely based outside of a jurisdiction are still legally accountable for the processing of data that originates in that jurisdiction and can be made subject to legal action in their own territory.

Data Localization

The concept of data localization is most relevant to multi-national organizations who operate in various jurisdictions and must deal with the challenge of storing and processing data where it is first collected and to persistently maintain its resident and sovereign status, effectively creating compute and data silos along jurisdictional boundaries. Multinationals will typically build business applications and deploy instances of them into each jurisdiction so that data and processing always remain local to the jurisdiction. Modern application architectures have made it easier to build applications that can span multiple countries in a distributed model. Such application architectures can maintain secure segregation of local data and processing with respect to each region while realizing the benefits and efficiencies of operating a single common application as opposed to maintain multiple instances globally.

Data Availability & Integrity

Protecting sovereign data not only involves preventing unauthorized or unintended exposing of the data to non-sovereign parties but also to make it highly available and ensure its integrity.

Consider the following recommendations for ensuring the availability and integrity of data:

- Always host workloads that handle sovereign data on highly available infrastructure that is capable of automatically restarting those workloads to a powered-on state in the event of a host hardware failure or power outage
- Implement high availability capabilities wherever possible on all components in the infrastructure that are involved in making data accessible, this includes but is not limited to network switches, NSX-T edge nodes, VSAN RAID levels, etc.
- Implement infrastructure redundancy in accordance with the design recommendations of software and hardware vendors
- Ensure that there are sufficient hosts in each cluster in the virtual infrastructure to tolerate host failures. In highly secure or remote environments in which new hardware lead times are extended, consider including more than one additional host for failure capacity purposes so that workloads can continue to operate as normal despite reduced overall hardware availability for long periods of time
- Design the infrastructure to include redundant links for both storage and networking, this may require additional redundant components in the host hardware
- For critical applications that lack native availability and redundancy capabilities, and which are of national importance, consider implementing VMware Fault Tolerance (FT). To ensure its effectiveness, enable high availability for all virtual machines involved in the application and its dependencies using a combination of VMware HA, FT, and application-level high availability so that no part of the protected application lacks redundancy
- Offer customers both virtual machine-level and application-aware backup and restore functionality. Ensure that these backup and restore capabilities are available in each security domain in the Sovereign Cloud, taking into account data classification isolation and segregation requirements
- Consider offering customers immutable storage that makes stored data completely tamper-proof, this should be offered with redundancy using either in-built replication or by offering a secondary instance

Sovereign Cloud Framework

This technical whitepaper discusses a subset of the Sovereign Cloud Framework for the design of a VMware Sovereign Cloud that guides partners through a self-assessment of market opportunity and applicable legal and security compliance frameworks to arrive at a conceptual, physical, and logical architecture design.

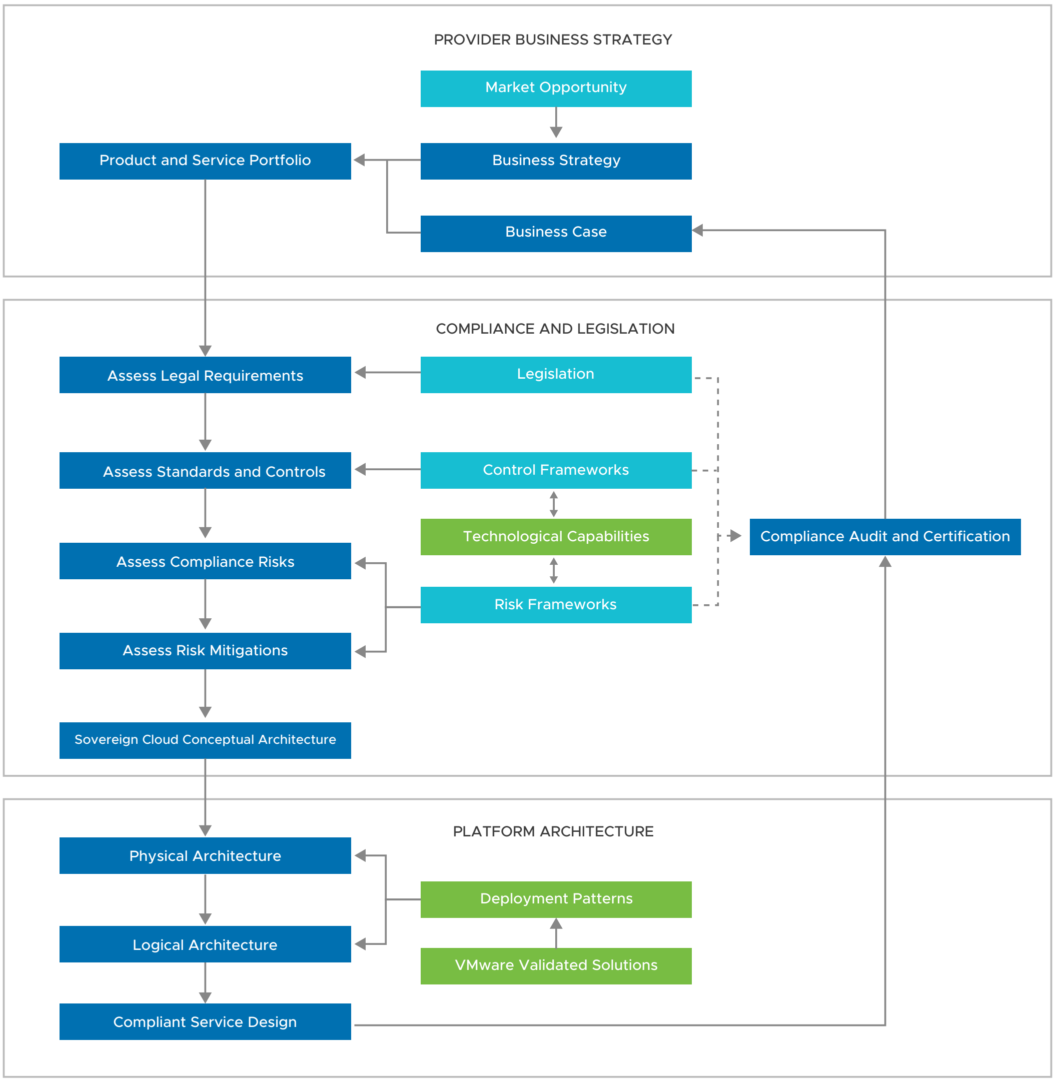
PROVIDER BUSINESS STRATEGY
Market Opportunity
Product and Service Portfolio
Business Strategy
Business Case
COMPLIANCE AND LEGISLATION
Assess Legal Requirements
Legislation
Assess Standards and Controls
Control Frameworks
Technological Capabilities
Compliance Audit and Certification
Assess Compliance Risks
Risk Frameworks
Assess Risk Mitigations
Sovereign Cloud Conceptual Architecture
PLATFORM ARCHITECTURE
Physical Architecture
Deployment Patterns
Logical Architecture
VMware Validated Solutions
Compliant Service Design

Figure 7 - VMware Sovereign Architecture Design Framework

The design framework is organised into three parts: Provider Business Strategy, Compliance & Legislation, and Platform Architecture. It represents a high-level process flow that starts with assessing business considerations and progresses to a phase of compliance self-assessment that defines the rules of operating the Sovereign Cloud in the provider's jurisdiction. The next phase involves defining an architecture that expresses how the Sovereign Cloud and compliant services will be offered technically. The final step is Compliance Audit and Certification which involves assessing the proposed architecture design against in-scope security frameworks to ensure compliance. The loop is then closed at the Business Case step to ensure ongoing and iterative development and maturity of the design, which is evaluated in the final stage in terms of a business case. Note that detailed design, build, and operationalisation are omitted from the architecture design framework for the purpose of brevity in this document. For more detail on these aspects please contact your VMware account manager.

This high-level design framework is offered as a starting point that can be further complemented by other frameworks if needed. It is acknowledged that partners who are considering developing a Sovereign Cloud offering may have already developed mature operating strategies for managing secure IT environments, possessing frameworks and methodologies for designing secure IT environments. Adopting the framework described here is not a requirement of a VMware Sovereign Cloud.

Provider Business Strategy

Before setting out on designing a VMware Sovereign Cloud platform it is necessary to identify the market opportunity and develop a portfolio of compelling service offerings that the provider can use to capitalise on and generate business. The service definitions provide architects with a scope and a mandate for developing a technical design and to identify technologies and capabilities that are required in the solution.

- **Market Opportunity –** A Sovereign Cloud seeks to address the challenges that enterprises increasingly face regarding data privacy and security laws. This is set against a backdrop of a continued increase in the use of public cloud, a trend set to continue, that has introduced additional technological and operational complexities as enterprises seek to remain compliant with the law. It is apparent that after transitioning from on-premises to a hybrid and multi-cloud strategy that many enterprises are now concerned that they are increasingly exposed to the risk of non-compliance and associated adverse legal and financial consequences. This increased awareness among enterprises of compliance risks highlights a future demand for Sovereign Cloud offerings.

  When evaluating the market opportunity as identified by the business, consider the following drivers impacting new and existing customers:

  - **Data Repatriation** – A need to return data back to sovereign territory to ensure compliancy, the data will likely be high in volume and hosted in a proprietary database platform in a hyperscaler today. As well as using like-for-like proprietary solutions, customers may consider alternative platforms for strategic reasons, such as open-source
  - **Sources or Demand for Value-add Services**
    - Services that customers want to support compliance with regulation
    - Services that customers use in hyperscalers and now need an alternative
    - Services that existing customers are asking for

The following table suggests some services that are of relevance and that would all be provided using resources entirely within the jurisdiction and in full compliance with data sovereignty legislation:

Service	Description
Virtual Machine and Container Workload Migration	Move virtual machine workloads from one platform resource to another within the Sovereign Cloud or across secure IT infrastructure either on-premises or on other cloud platforms while maintaining compliancy
Database Migration	Assisted database migration services with alignment to data classification and compliance principles
Application Migration	Assisted application migration services with alignment to Sovereign Cloud security domain architecture patterns and security compliance
Assisting Sovereign On-boarding	Offer specialist assistance for importing workloads, data, and applications into the Sovereign Cloud in the most compliant manner. Give customers guidance on strategies for making best use of the security and compliance features of the platform so that customer applications, data and connectivity properly align with their compliance objectives
Managed Services	Provide management and specialised engineering and consultancy services to newly on-boarded customers, either as an optional or mandatory service component
DR	Protect customer workloads and data so that recovery is possible in the event of a disaster
Secure Connectivity	Enable secure and compliant connectivity services for customers between sovereign endpoints, on-premises infrastructure
Database / Data Warehouse / Data Lake Services / Data Analytics	Customers who are considering repatriating their data to their 'home' jurisdiction may opt to import their data into a new data infrastructure rather than migrating database virtual machines onto the platform. These customers may be using a Data Warehouse or Data Lake services on a hyperscaler already. Customers may have very large data sets not suited to traditional single or clustered database instances or may be using a DBaaS platform service that does not allow whole-VM export
AI & Machine-Learning	Model training using machine-learning and deep-learning capabilities powered by CPU, GPU and FPGA, with processing proximate to the data being analysed (integrated with data infrastructure services hosted on the Sovereign Cloud Platform itself)

For a broader discussion on the case for Sovereign Cloud and the opportunities it represents for providers and enterprises alike, please refer to VMware's Sovereign Cloud Commercial Whitepaper.

- **Business Strategy –** The business strategy of the provider outlines the plan that must be executed for the business to meet its goals. Such goals will include how to go about addressing identified market opportunities, how to grow the business both in terms of the number of customers and in the level of business the provider does with each customer on an ongoing basis. The business strategy may involve targeting new and existing customers from specific industries and sectors, these may represent new industries for the provider that could potentially require changes to business processes, hiring new personnel to expand skillsets and acquiring certifications, licences, and accreditations. If the provider must make such changes, then the architect should work closely with all relevant stakeholders on an ongoing basis to ensure consistency between the business and the technology.

  An architecture design must clearly describe how it will support the business strategy, from a technical standpoint this means ensuring that it includes functionality that will serve the demands of customers and that the provider needs to offer services that are relevant. It also involves ensuring that the Sovereign Cloud platform is appropriately scaled in terms of capacity and performance for anticipated demand from the very beginning and can scale appropriately with forecasted demand.

- **Product & Service Portfolio –** A portfolio of service offerings is developed that together are designed to appeal to prospective customers within the markets identified in the business strategy. This activity typically falls under the ownership of a portfolio or service manager. A service definition is produced that outlines the functionality, features, and characteristics of the service and any applicable upper and lower bounds, flavors, etc.

  The primary purpose of a Sovereign Cloud is to provide those enterprises with data sovereignty and privacy concerns with an assisted and accelerated path to self-directed compliancy and assurance, ideally at a competitive price point. By incorporating many of the desired security, governance, and audit characteristics into the platform that can address these challenges the provider can create offerings that enterprises find compelling. Some of these characteristics are not primarily technological. Simply providing a platform that is geographically local and which is owned and operated by a local legal entity whose ultimate majority ownership is also local is often not enough. The technical aspect is concerned with ensuring that data does not lose its designated protective status at any point.

  The range of services in the portfolio should be broadly like that of a commercial cloud offering but with some enhanced security provisions and processes integrated into the product. Therefore, the sovereign status of the platform should not represent a barrier to developing value-add services for customers but rather the opposite. A simple example of this principle is DR services. This can be offered using the same tools and processes as a commercial cloud except the failover destination must also reside within the same geographical jurisdiction. No element of the service can be present in or managed by non-resident entities. The service is encapsulated by security controls that prevent unintended access to data that have the potential to lose sovereign protected status.

The Sovereign Cloud Framework takes a layered approach to service and architecture design in which the key requirements and characteristics of industry, sovereignty, and residency are addressed independently but are developed to serve one another.

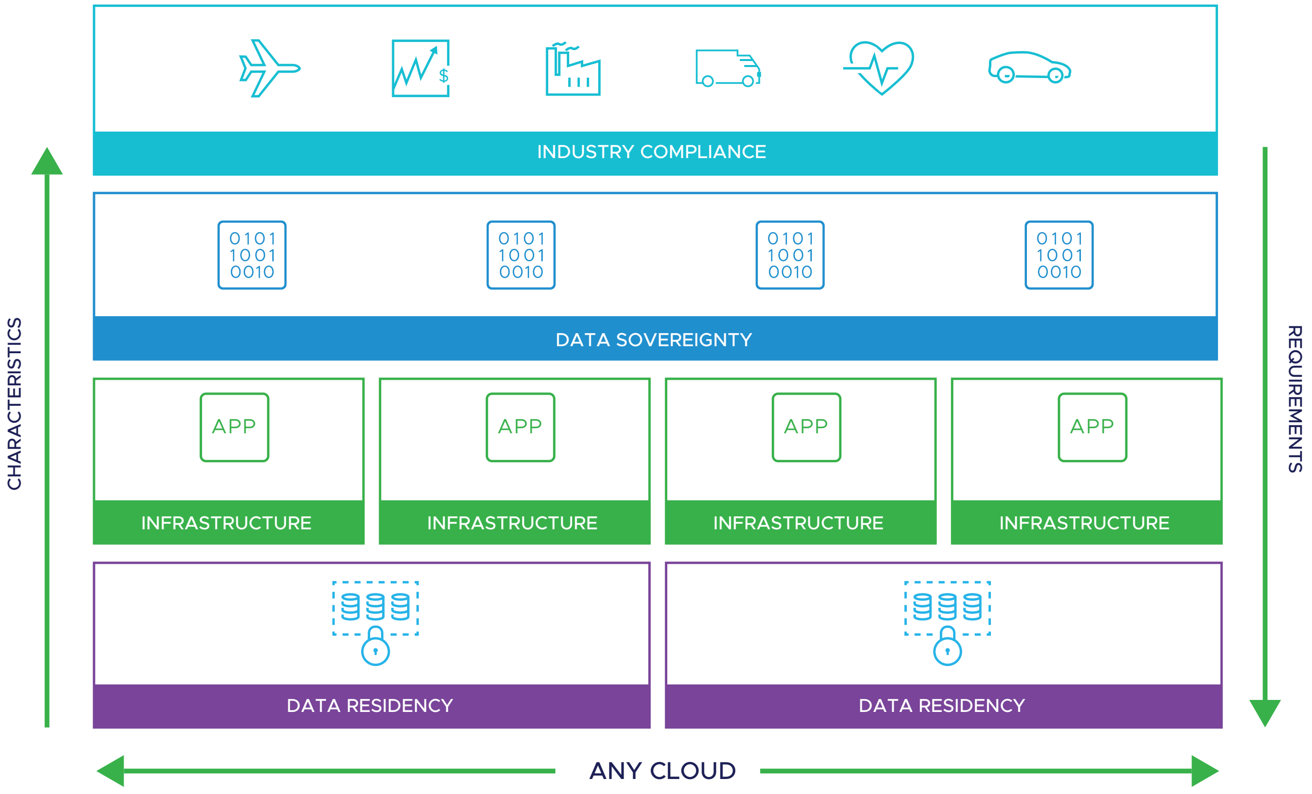

Figure 8 - Layered Approach to Compliant Architecture Design

Points:

- Requirements specific to each layer are handled independently. The service and technical characteristics of each layer drive requirements down to the next layer and utilize the characteristics presented by the layers beneath
- The Data Residency layer addresses the need to restrict the physical location of data stored at-rest to a specific country or region. Its function is to ensure that the data never resides outside of that boundary. This is addressed by locating physical infrastructure within the jurisdiction and using classification and categorization for the purpose of placement and policy-based automation
- Infrastructure characteristics such as availability, resilience, redundancy, and backup & restore are implemented at the data residency and infrastructure layers
- The Data Sovereignty layer ensures that no matter where the data is stored or processed it always remains under the sovereign control of the owning jurisdiction. This allows for the opportunity for controlled distribution and processing of data beyond the resident boundary under controlled conditions
- The Industry Compliance layer represents the boundary of enforcement for conventional industry-specific compliancy standard and controls such as PCI, HIPAA, etc. This layer drives many requirements for specific service and technical characteristics necessary to achieve compliance down through the underlying layers

- **Business Case –** The development of a business case for a new or improved service is an activity that starts immediately once the service is proposed and identified as a candidate for market launch. It requires an analysis of the estimated cost of design, implementation, and operation throughout the projected lifecycle of the service. The cost is then evaluated against the proposed pricing model and forecasted demand to ultimately arrive at an understanding of the probable returns on the investment the service represents. To undertake the cost analysis, the architect must propose a bill of materials that encompasses the hardware and software needed to build the service at forecasted scales. This information provides business stakeholders with an understanding of the required initial capital expenditure as well as an indication of likely operational costs related to software licensing, subscriptions, support contracts, and staff. Ideally, the scope of the business case will be sufficiently wide to cover a range of services and at sufficient scale to allow for a true appreciation of cost efficiencies at-scale. A key consideration behind VMware Sovereign Cloud is to offer providers with opportunities to drive down the cost of ownership as much as is reasonably possible using efficiencies in technology so that offerings can be priced competitively.

  The architect should consider the range of services the business is considering offering and the scale at which they are forecasted to be delivered.

Compliance & Legislation

In this phase of the design framework the legalization and security compliance frameworks that are applicable to the services being proposed are identified and their technical implementations considered. The conclusion of this phase is a conceptual design that reflects the security classifications and security domains that the platform must include to be appropriate to the jurisdiction and market segments being targeted.

- **Assess Legal Requirements –** The first order of business when determining the rules for handling data is to understand the legal obligations that apply to the provider in the jurisdiction when hosting customer data. Legislation relating to data security and privacy is well developed in most nations and very often the laws of the land place responsibilities with those who store and process data while also assigning rights to data owners and subjects that data processors and owners are obliged to protect. Furthermore, aspects of data privacy are enforced in international law, which means that legal rights and obligations can even extend to parties outside of the jurisdiction that is local to the provider's operations.

  It will very often be the case that there will be multiple laws and statutes that relate to data and privacy, in some cases there may be laws related to specific regulated industries such as finance, healthcare, nuclear infrastructure, etc. Seek legal advice from a qualified legal practitioner in the applicable jurisdictions to obtain a full briefing on what laws apply, including those relating to specific industries.

  To be clear, in many parts of the world the cloud provider shares in the legal responsibility for data handling and processing with the customer. There is a potential risk for legal exposure in the event of a breach depending on the circumstances as well as obligations to act when another party wishes to exercise their legal rights over data.

  In order for the service provider to mitigate against legal risk, the architect should determine what security features and functionality should be enforced in the service offering appropriate to the security classification, e.g., forced encryption of data at rest for customer workloads with no ability for the customer to opt-out. An additional consideration is ensuring that end-customer data owners and controllers are provided with tools and mechanisms for exercising their legal rights, either technologically or by way of processes supported by the provider.

- **Assess Standards & Controls –** The provider must determine which standards and control frameworks, such as NIST, it should adhere to, this may be dictated to some degree by local legislation and the customers it will service, such as Government and Military. A range of security frameworks stipulate controls that must be implemented in secure computing environments and a set of minimal standards to adhere to.

  The architect should catalog all applicable controls and standards that must be applied to later apply them in the design.

- **Assess Compliance Risks -** In addition to Standards and Control frameworks, other frameworks take a risk-based view of IT security to identify various risks related to technology and process. Then offer mitigation recommendations.

  Consider the following additional risks that may apply and determine how best to mitigate them:

  - Data security legislation becomes more restrictive or complex in the future, not less
  - Products from specific foreign hardware and software vendors may be prohibited for use in cloud hosting environments. Similarly, a negative perception of vendor products from certain foreign states may impact a customer's willingness to consume services from a provider even if the use of such products is lawful. Possible mitigations include avoiding the use of products that originate from foreign states that have less favourable relationships with the home jurisdiction in geopolitical terms. Try to select vendors from within the same jurisdiction or from other jurisdictions with whom the home state has longstanding and publicly acknowledged positive geopolitical relationships as well as similar or even reciprocal consumer and business law. If this is not feasible then consider using other technologies to ensure that applicable devices never process or store unencrypted data.
  - The financial position of the provider changes and becomes insolvent. What happens to the data if the provider collapses?
  - Personnel have been legitimately granted access to the platform as part of their role, however, they have been improperly unvetted in accordance with local employment or security rules
  - Hosted data has been incorrectly classified and as such is not protected in a compliant manner
  - Sovereign data becomes accessible to systems either within or outside of the platform that are not designated sovereign and may reside outside of the jurisdiction without suitable controls in-place to maintain the sovereign status of the data

  These are just a few examples of possible risks. It is recommended that a thorough risk assessment is conducted that identifies the risks that, if realized, result in a loss of sovereign control of data hosted in the platform.

- **Assess Risk Mitigations –** Having assessed the applicable standards, controls and risks that could apply to a Sovereign Cloud, the architect should identify all the security capabilities and features they have available within their chosen product sets and match them against the security controls that are needed. If any controls appear to be lacking, then the architect should determine what additional products might be needed to introduce the desired controls or consider custom orchestration or manual process as alternative solutions.

  VMware provides a comprehensive register to compliance controls for various frameworks in VMware products and highlights configuration options. These are most useful at the detailed design phase which is not covered by this white paper but provides a useful insight to determine the level of coverage available for the controls that are required in the platform. The link is provided **here**

- **Sovereign Cloud Conceptual Architecture -** Before a technical architecture can be defined for a VMware Sovereign Cloud platform it is first necessary to define a conceptual model of the security and data classification environment. The goal of the conceptual architecture is to model the security domains and classifications that exist, describe their rules for data and workload placement, interconnectivity between the security domains and intra-connectivity within them.

  – **Resident & Sovereign Domains** - VMware Sovereign Cloud stipulates a minimum of two enforceable security domains in security model as a minimum.

  – **The Outside World** – Any entity that exists outside of the Resident and Sovereign domains may be considered the outside world, this extends to other systems and people in the same service provider business that are not designated Sovereign as well as the public internet. Essentially, this domain represents anything where data sovereignty and privacy cannot be enforced with a high degree of confidence.

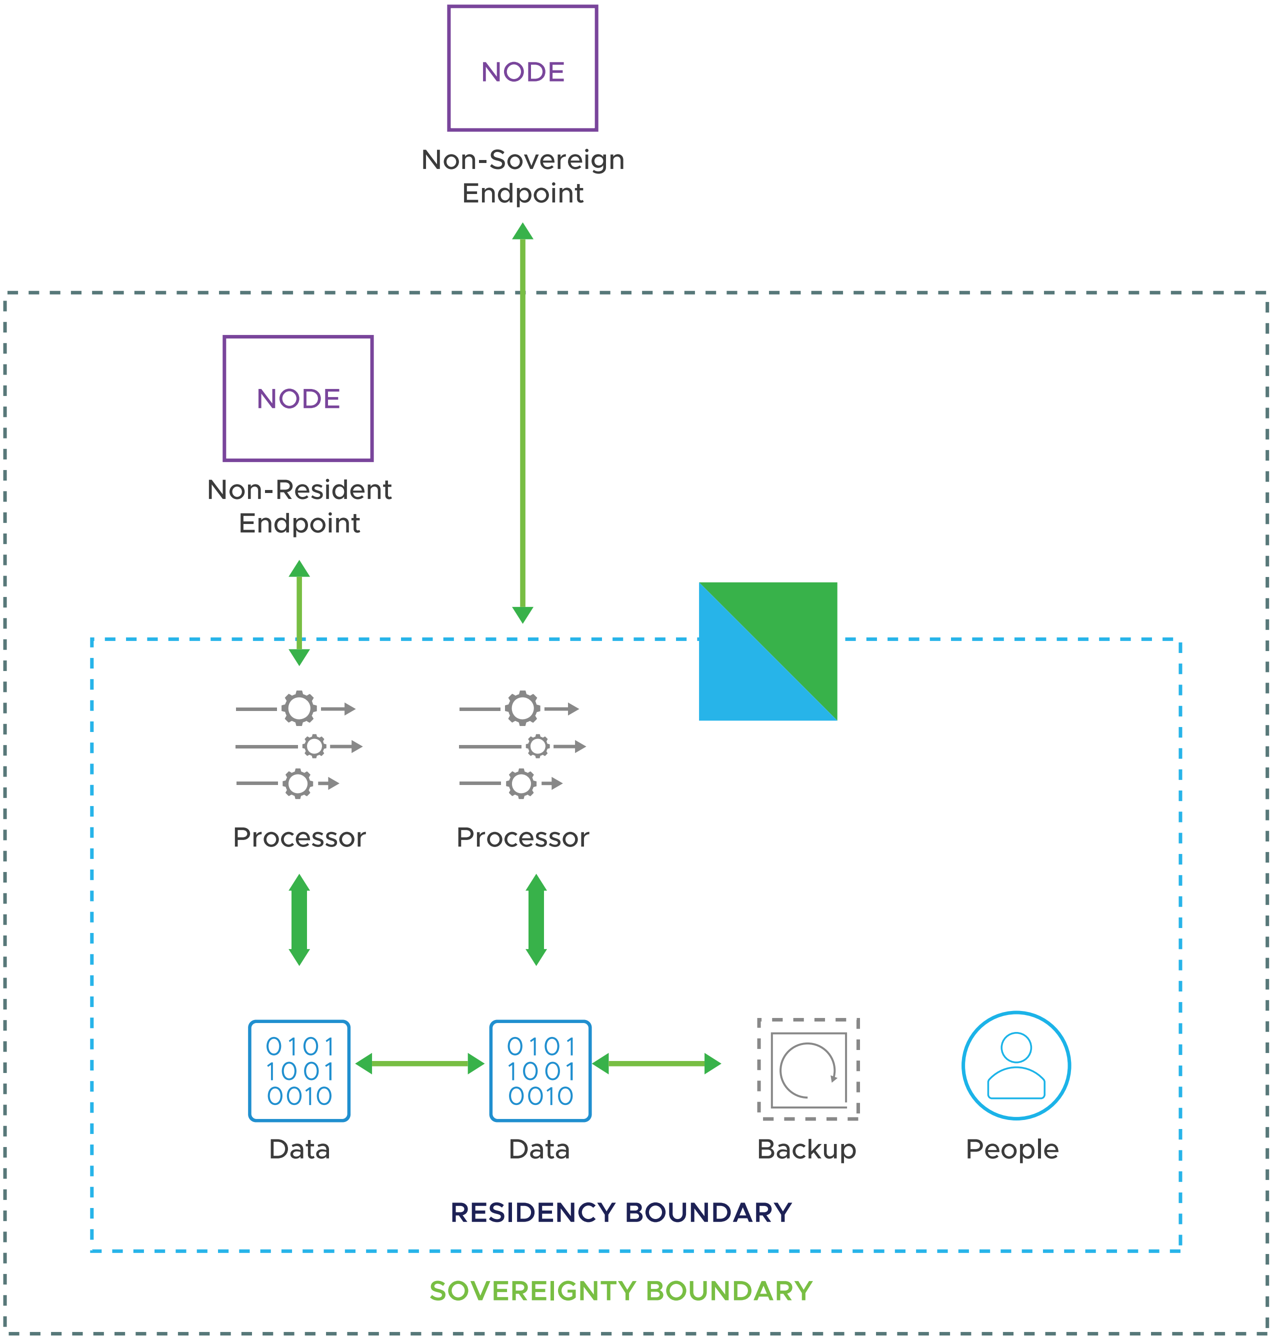

Figure 9 - Cross-domain Interactions

- **Security Domains for Data Classifications** – Identify recognized official data and system security classifications appropriate to the jurisdiction. Examples include official Government classifications such as Confidential, Secret, Top Secret, NOFORN, etc. Examples from the commercial sector can include Confidential, Internal Use, Public, Sensitive, Highly Sensitive. The classifications that the provider chooses to include in the platform by default will depend on a combination of local jurisdictional norms and the type of customers the platform is intended to serve. In some cases, it may be sufficient to simply offer the Resident and Sovereign domain classifications and allow customers to manage their data classifications independently of the provider.

The handling of data at various security levels is normally subject to strict rules on how the data is stored, processed, and secured. It is unusual for systems from different classifications to interact and for a system to host data sets with differing classifications. Any data record, metadata, or system within the bounds of a location, system, or storage medium with a specific security classification assigned will assume that security classification automatically. This principle is reflected in a VMware Sovereign Cloud in technology and policy-driven management and security. Moving data from one classification to another is normally a controlled process as it has the potential to have significant consequences. For this reason, the ability to move data between classifications should be heavily restricted. Movement of data between classifications as well as access to data with certain security classifications will routinely require approval from a suitable authority to prevent data leaks which could have legal consequences. In many countries, the security of data classification can change based on factors such as the number of records held or the nature of any one particular record which can elevate the classification of the entire data set. Lowering the classification level of data is even more unusual and is essentially not done. There is also the risk of contamination leading to service interruption, if a secret record accidently enters a non-secret data set, then the entire data set is now considered secret. This would result in the data set having to be moved to the secret infrastructure and for the non-secret infrastructure to potentially be securely wiped, both actions would result in service outage.

The responsibility for handling data and assigning classifications is not usually the remit of the provider, however, once the provider is responsible for handling securely classified data from the Government and Commercial sectors then they are also likely to be bound to the rules that accompany those classifications. This means that approval for certain platform management operations may have to be sought from customers and security officers.

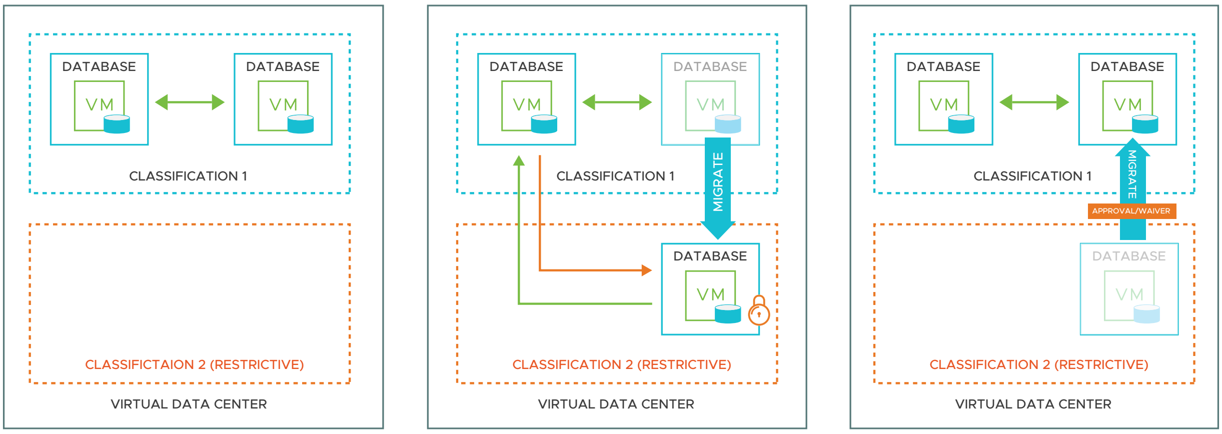

Data within a classification can move and connect freely.

Migration or replication of data to more restrictive classification causes inheritance of restrictive characteristics. Storage becomes encrypted. Communication rules do not allow inbound connections into classification 2, but outbound connection is allowed.

Migration or replication of data back to less restrictive classification prevented by policy. This is only possible on approval or waiver of the authorized party to certify the data can leave its protected status.

Figure 10 - Workload Security Classification Transitions

In Figure 10, two classification levels for data are depicted that have different security policies. Systems within classification 1 may freely connect with each other and their storage does not have to be encrypted. Systems within classification 2 may establish connections to systems in classification 1 to query data, but this can only happen in one direction; connections from classification 1 to classification 2 are not possible. Workload migration may take place from classification 1 to 2 without explicit approval as the target classification has a more restrictive security posture. As the system is migrated from classification 1 to 2, it inherits the encrypted storage policy of the target classification, and the storage becomes encrypted to be compliant. The migrated workload that was once reachable by the other system in classification 1 can no longer be reached. Finally, to return the migrated system from classification 2 to 1, the migration can only be completed if approval is explicitly granted. This is necessary as data that resides in classification 2 must be treated in a more restrictive manner in this example and may not necessarily be allowed to be moved to a less restrictive classification.

- **Define Security Rules for Security Domains** – Based on the information gathered and security decisions made in the Compliance & Legislation phase of the framework, it should be possible to define clear rules for the placement and interaction of data and systems. In VMware platforms, these rules are reflected in administrative organisation of systems and software defined networks. They are also reflected in network security policy, firewall and micro-segmentation policy, storage policy, user roles and approval mechanisms.

  The output of this step in the process is required in the detailed design phase.

When fully expanded, the VMware Sovereign Cloud architecture from the perspective of the tenant might look like this:

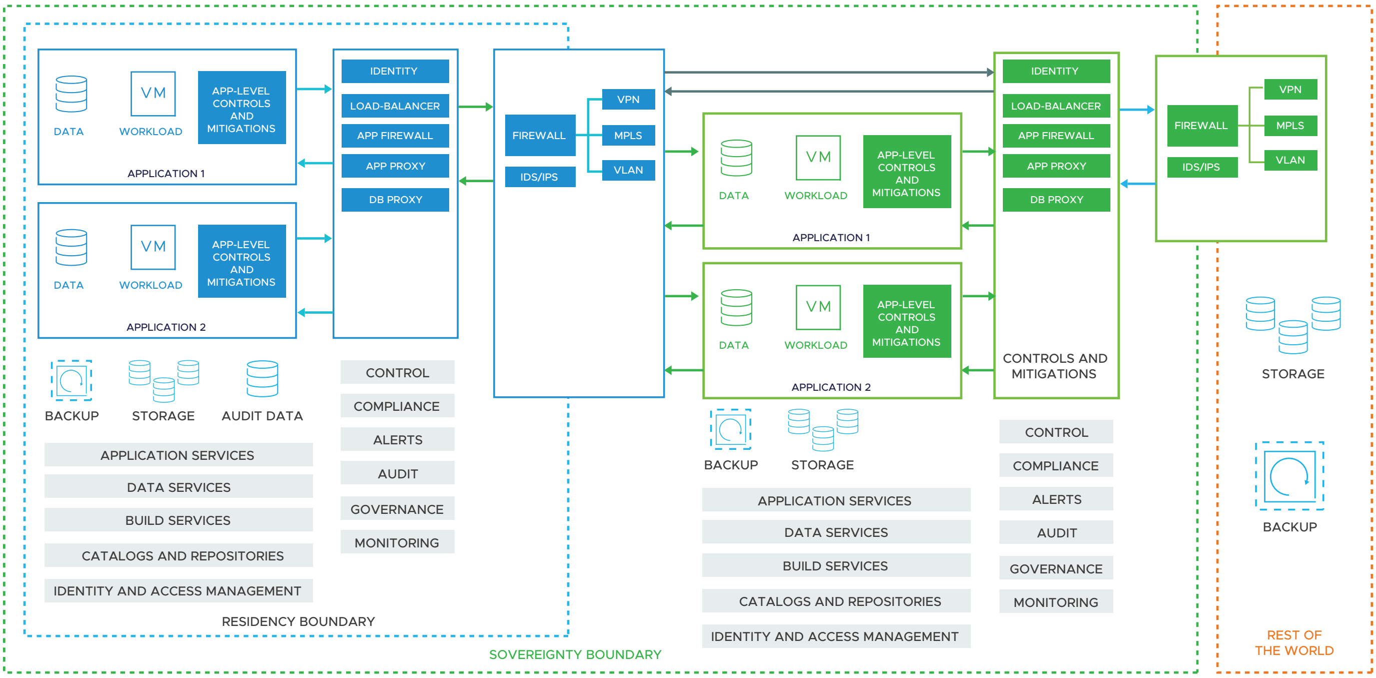

Figure 11 - Sovereign Cloud Detailed Architecture Pattern from a Tenant Perspective

In the above illustration, two applications are depicted, and both have a presence in the Resident domain (data tier in blue) and the Sovereign domain (app tier in green). Communication between the application tiers and the outside world passes through the domain boundary controls and mitigations which are composed of all the security and control services that the provider is offering. The customer also could insert their own additional security controls at the application level. The applications depend on backup, storage, and other services and capabilities offered by the provider. These are offered separately within each domain to ensure secure isolation of data flow and the avoidance of unintended metadata exposure. Audit data is gathered from both domains but is stored exclusively in the Resident domain.

- **Compliance Audit & Certification –** Once the architect has prepared a draft design of the solution it should then be re-assessed independently against the original security criteria gathered in the Compliance and Legislation phase. Depending on the legal requirements or customer requirements in the jurisdiction it may be necessary to seek an independent audit from external experts who can certify or accredit the design for use.

  Output from this process should be made available to prospective customers on a confidential basis to demonstrate the secure status and credentials of the Sovereign Cloud platform, with regular audits and reviews made available to existing customers of the platform, ideally on a continuous basis.

Platform Architecture

The key objectives of a VMware Sovereign Cloud platform architecture are:

- Provide a highly secure computing platform
- Provide intrinsic end-to-end security consistently across the solution
- Support enforcement of data sovereignty and residency boundaries
- Support the segregation of data classifications
- Suitable for both secure multi-tenancy and customer-dedicated deployment models
- Suitable for Commercial, Industry and Government customers
- Optional controlled and secured connectivity beyond the platform

The basis of the VMware Sovereign Cloud platform architecture is the VMware Validated Design for SDDC (VVD) which is a publicly available reference architecture for the implementation of a VMware SDDC which has been validated for component interoperability, scalability, and supportability. It describes a base architecture for the deployment of VMware SDDC components in relation to physical sites and compute domains aligned to consistent networking and storage. The VMware Sovereign Cloud Architecture is represented as an overlay of design patterns and constraints to the VVD that together deliver secure compute services in a flexible way.

VMware Cloud Foundation (VCF) offers an accelerated way to design and build supportable instances of a VMware SDDC using a combination of a curated configuration definition files that engineers pre-populate, validate and automate to streamline a SDDC deployment against the configuration provided. The automation extends to lifecycle operations of the platform which simplifies the process of adding additional hosts and clusters and applying patches and updates. The deployment process of the platform can be further automated with additional tools such as PowerCLI and vRealize Orchestrator.

VCF offers service providers an opportunity to rapidly build secure virtualization platforms on available supported hardware to reduce overall provisioning time to end-customers. The principle of a highly repeatable and automated platform deployment allows for the possibility of production-line style rapid platform building that can greatly reduce lead times when introducing new capacity.

In the case of a Sovereign Cloud, automated platform provisioning can be an attractive capability when working with select Government and Military customers who require provider-built and validated platforms that are then operated at-distance (such as during operational deployment) or that are not accessible to provider engineers and management tools once delivered. Similarly, in highly sensitive environments where the lifecycle of the platform is drastically reduced from years to weeks or months, the platform is ‘reset’ at the end of the operation before being rapidly repurposed for another operation. Secure compute platforms that are operated at-distance and where remote access is either limited or unfeasible such as on naval deployments, battlefield deployments or overseas foreign embassies, must operate reliably and predictably. Using a combination of standardised reference architectures and highly repeatable provisioning is highly recommended in these cases.

Following the development of a conceptual security architecture for the platform, a more traditional technical design methodology can be adopted to develop the physical and logical architecture designs. Using VMware’s Validated Designs or Validated Solutions, a supportable VMware SDDC or VCF can be designed as a base platform for Sovereign Cloud services to be overlaid.

- **Physical Architecture –** The objective here is to specify the physical structure of the platform necessary to implement the required number of host clusters organized into SDDC domains, while supporting the physical separation requirements of the conceptual security architecture. Host clusters represented in both Management and Workload domains in a VMware SDDC are aligned to network switches, classifications, physical sites, etc. Additional physical security elements such as secure server cages for the placement of hosts and switches, or fire-proof server enclosures should also be referenced.

The physical architecture of a VMware platform is generally aligned to sites and the segregation and distribution of SDDC domain clusters both within and across sites. A VMware Sovereign Cloud introduces additional considerations brought about by security classifications. As an example, Government Secret and Top-Secret systems are physically separated from those that handle lower classifications or those that are unclassified.

Depending on the security classification regime in the local jurisdiction, it may be possible to group hardware for some security classifications together. To reduce overall cost, it is recommended to favor logical separation over physical separation where possible so long as security compliance can be maintained. There will sometimes be exceptions, some customers may insist on a physical segregation (which would also extend to the management plane). This should be articulated in the architecture design as one possible deployment for the purpose of technical validation but may not necessarily be marketed as the default offering.

Note that the VMware Sovereign Cloud Framework does not stipulate physical segregation of the Resident and Sovereign security domains, however separate SDDC domains and therefore workload clusters for each is recommended, this results in a partially shared management plane. The provider may wish to take this further and have fully dedicated SDDCs for each security domain which is also a valid option.

For providers who are seeking to target Government and Military customers or private sector customers looking to handle classified data then it is recommended that the architecture proposes two physical SDDCs that can be physically isolated from one-another. One physical SDDC would be dedicated to highly classified data and the other be dedicated to data of lower or no security classification. The implication is that data with a low security classification can reside within the same physical platform as unclassified data with suitable security controls and separation in-place. This overall pattern could be suitable for classified secure environments in many Western nations and may represent a good starting point.

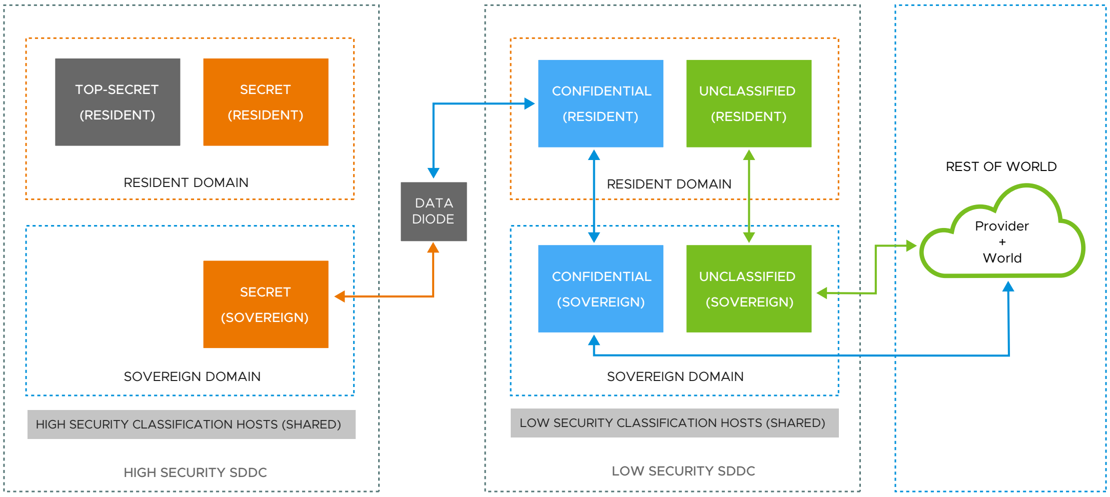

Figure 12 - Physical vs Logical Security Separation Example

In the above illustration, the two SDDC stacks are shown with their respective Resident and Sovereign security domains. This is achieved by deploying two VMware Sovereign Cloud architecture patterns side-by-side, optionally using VCF. As network traffic cannot reach Resident domains from the outside world it is necessary to present Sovereign domains for each classification that can expect connections from outside. In this example, top-secret workloads never connect with non-resident, non-secret, or non-sovereign systems so no Top-Secret Sovereign domain is necessary. Connections can be established from the low security SDDC to the high security SDDC via a data diode which in this example is controlled by the customer and which prevents secret and top-secret data from leaving the environment. Again, data flowing via the data diode can only be presented to the Secret Sovereign domain as direct Resident connectivity is not possible. It would be entirely possible for Top-Secret workloads to communicate with Secret workloads within the same Resident domain if both parties allowed this.

In the above example, physical separation of the various domains is possible but is not necessary to satisfy security requirements, therefore relying on a mix of physical and logical separation. It could also be the case that each security domain consists of multiple tenants as they are segregated by other means.

From a management plane perspective, SDDC management workloads are hosted in dedicated management clusters in a non-consolidated scenario. This includes the vCenters and NSX-T Managers for each SDDC workload domain any additional infrastructure workloads such as Backup servers. This means that from a security classification perspective, the management plane of the entire SDDC instance inherits the highest security classification of the stack. Provider engineers operating the high security SDDC instance would have to be cleared to interact with top-secret systems even if measures to restrict access are taken. As there are two SDDC instances in the previous example, this presents two Management SDDC domains and two physical management clusters that are not presented in this diagram but are omitted.

- **Logical Architecture –** The logical architecture of the platform relates to the functional components of a VMware SDDC and how they will interact with each other as well as with external infrastructure dependencies such as DNS, authentication services, etc. VMware Validated Designs, VMware Validated Solutions and VCF all describe supportable reference logical architectures for a VMware SDDC. In a Sovereign Cloud context, the physical architecture addresses how host clusters are organized to address the physical separation requirements of the security domains specified in the conceptual security architecture. Where logical separation is stipulated, such as for shared host clusters for workloads from different data security classifications, then the logical architecture describes how they are organized using VMware logical constructs. These can include DRS resource pools, virtual infrastructure inventory folders, NSX security groups, transport zones, etc.

  For example, a workload cluster can include a separate DRS resource pool and inventory folder for each classification. This creates both a resource and administrative boundary around each data classification that can be used for access control and resource management purposes. Each DRS resource pool can be represented in VMware Cloud Director as a provider vDC for a given classification. Allowing for classification-specific Org vDCs to be offered to tenants and offering the additional benefit of logical network segregation. Tenants determine which Org vDC to provision their workloads to, depending on their designated data classifications and can create network groups and rules to manage those classifications individually. The logical architecture should describe how security domains and classifications will be represented in the SDDC to customers and how workload placement and policy-based management strategies can be implemented by the provider to consistently enforce compliance.

- **Compliant Service Design –** The principle of a compliant service design is to ensure that services are developed by the provider to cater to each security domain and classification that the platform will service. For example, if the provider is offering services to Government, then there should be a backup service, however, due to there being at least two security domains there must be an offering for Resident Backup and Sovereign Backup. These backup services must be delivered entirely from within their respective security domains and will represent two sets of separate backup infrastructure that are isolated from one-another. The implication here is that these domain-specific backup services cannot be used for backing up a workload in one domain and restoring it in another. This principle may be extended further for data security classifications in which local compliance regulations require separation of management infrastructure between specific classifications.

  Compliant services will in most cases represent traditional IaaS, PaaS and even SaaS services found in a typical hosting environment but with additional characteristics and constraints related to security. For example, this could include selecting a classification status when provisioning a new virtual machine. The cluster in which the virtual machine is provisioned, and which networks can be attached to could depend on which classification it is assigned. This may also impact who may manage, access, and connect to that virtual machine once it is provisioned. Similarly, assigning a data classification or some other protected status could require the enforcement of certain functionality such as disk encryption, DR protection, or zero-trust micro-segmentation for the purpose of enforcing compliance. It may be the case that sovereign data in a production environment must always be replicated or backed up without exception. If this is the case then backup should not be offered as an option but as an integrated characteristic of the service. When developing service offerings, the service provider should consider what security and service characteristics should form part of a service and seek to present appropriate service options and eventually automate them where possible.

Sovereign Cloud Architecture

Sovereign Cloud Conceptual Architecture

The conceptual architecture for a Sovereign Cloud is composed of the technology components of a reference VMware Cloud Provider Platform overlaid on the conceptual compliancy framework.

The logical building blocks of the reference architecture are the management and workload domains of the VMware VVD reference architecture which is also the architectural basis for VMware Cloud Foundation. The Resident, Sovereign, and Data Classification security domain instances that are identified in the Conceptual Security Architecture are overlayed onto domains.

Tenant Network Perspective

The following illustration depicts one possible network architecture of a VMware Sovereign Cloud based on NSX-T and 3rd party physical network security appliances protecting the domain boundaries.

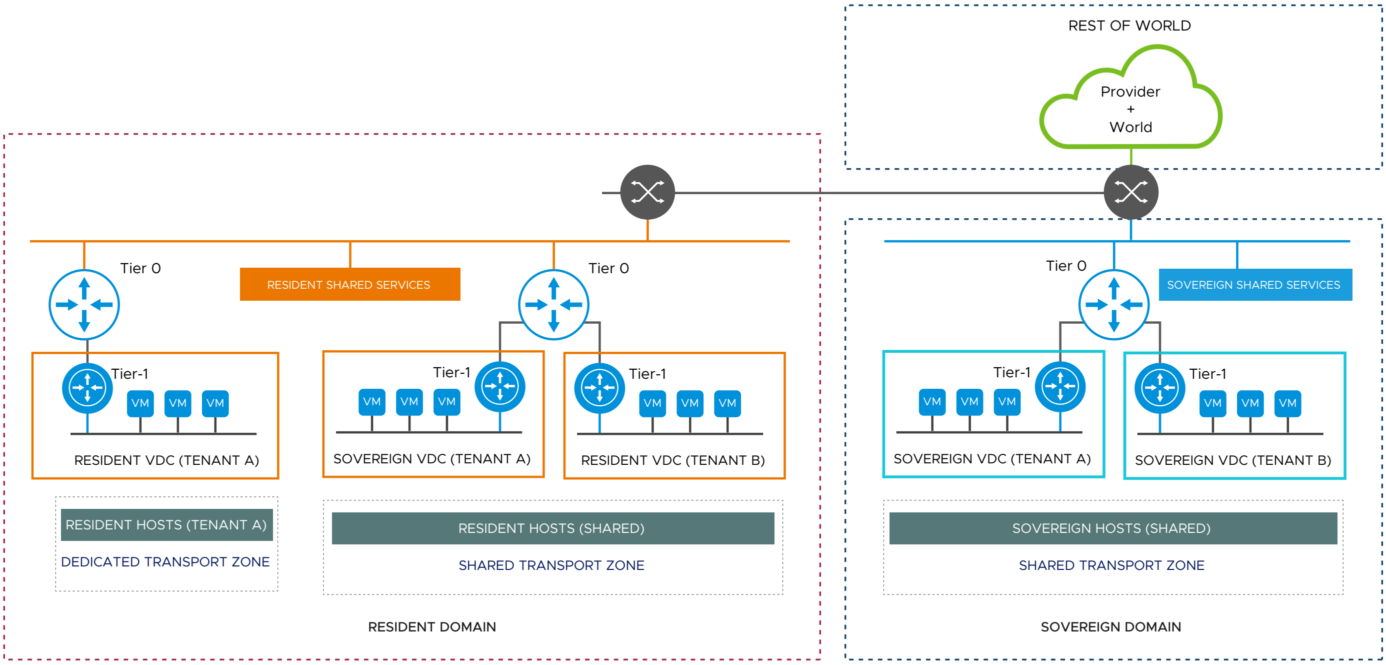

Figure 13 - Physical vs Logical Security Separation With 3rd Party Appliance Example

In this example, tenants A and B have both resident and sovereign resources provisioned. Tenant A has a combination of dedicated resources (physical hosts) as well as resources allocated from a shared pool. Shared services offered to all tenants by the provider such as backup, software update repositories or management and automation tools, are locally connected within each domain.

The resident domain is analogous to a garden walled network and storage connectivity, there is no default path for data to pass in or out. The sovereign domain is analogous to a typical DMZ network connectivity area as found in all internet-facing hosting environments, systems in the domain can interact with the outside world under controlled conditions and act as a front-end to an application where the data resides in the resident domain.

In the following illustration, a similar network architecture is depicted but in this case the topology is simplified further by depending solely on NSX-T Tier 0 gateways at the domain boundaries. This approach reduces the overall need for physical network security appliances at the domain boundaries but still allows for optional 3rd party network security integrations that support NSX-T service insertion that can be delivered using virtual appliances.

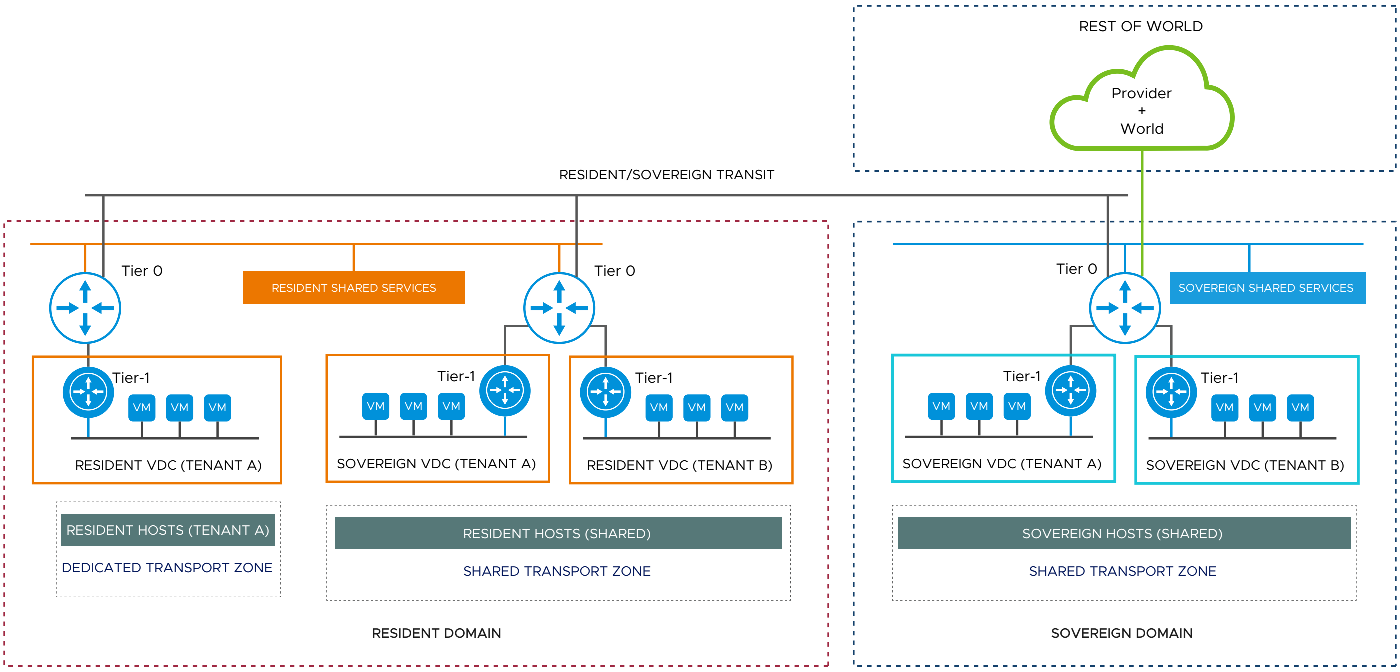

Figure 14 - Physical vs Logical Security Separation with NSX Only Example

Network connectivity from the Resident domain can only go outbound to the Sovereign domain. The Tier 0 gateways are configured with black hole default gateways and depend on static routes to enable connectivity to specific networks in the Sovereign domain and to other Resident IP ranges. Connectivity between the Resident and Sovereign domains is via the Resident/Sovereign Transit network which, if disconnected, makes all Resident virtual machines inaccessible to Sovereign systems. The Tier 0 gateway in the Sovereign domain does not advertise routes up-stream, depending entirely on static route configuration by the provider.

In this example, only NAT and Load-balancing is used when connecting between the Resident and Sovereign domain, no tenant network routing is allowed. This obscures the IP addresses of systems on either side.

The Compliancy Chain

The principle of the compliancy chain is to extend the same level of trust, security and sovereign assurance to customer-owned workloads and data in other platforms as can be applied to workloads hosted in a VMware Sovereign Cloud. The need for a compliancy chain is to acknowledge that most IT systems do not function in isolation and at some point, must interface with less trusted systems. Applications may even have dependencies on external systems where trust cannot be guaranteed.

The necessity for a compliancy chain is brought about by the fact that in most cases customer workloads in a Sovereign Cloud environment will be interconnected with other trusted systems. These can include customer on-premises systems, secure private network circuits, and even resources on other cloud platforms, including hyperscalers. While Military and Intelligence customers are usually an exception, most secure IT environments will at some point interact with other systems that are not entirely under the provider or customer's control, such as client devices. In a compliancy chain, customer data may be transferred and processed off-platform with suitable trust mechanisms and controls in-place. This allows customers of Sovereign Clouds to take advantage of many of the services that hyperscalers provide while keeping the most sensitive data and systems on home soil.

An example of this principle would be if a data set in the Sovereign Cloud were to be replicated to a hyperscaler platform in another country to be processed by an application using functions-as-a-service in the remote platform. In this example the data is encrypted at-rest or is kept entirely memory resident while in the hyperscaler using a product such as VMware Tanzu GemFire. The data replication traffic is also encrypted so that interception is not possible at any stage. The data itself might also be obfuscated or tokenized so that the true informational value of the data is never exposed to the outside world but is still processible in a useful way. The outputted 'product' of the processed data is immediately replicated back to the Sovereign Cloud so that the primary data set and any new entries always persist within the 'home' jurisdiction. Any asymmetric private encryption keys, certificate authorities, or identity and access management involved in the compliancy chain remain fully within the jurisdiction and ideally within the Sovereign Cloud, out of reach of foreign actors. No identity and access management functionality in the hyperscaler is used by the sovereign application which instead entirely depends on services hosted in the Sovereign Cloud, reachable via encrypted connectivity.

It is possible to maintain the sovereign status of the data set in this example even when working with the data sets across cloud platforms, including those that are outside of the jurisdiction. The sovereign data does not get exposed to non-sovereign entities, is revokable by the data owner, and its integrity is protected as it cannot be tampered with.

Cloud providers have an opportunity to develop services around the concept of the compliancy chain. This will be of significant benefit to enterprises with a multi-cloud strategy. Using products such as VMware vRA, NSX and vRO makes it possible to provision, manage and connect workloads and services on both VMware and native hyperscaler platforms and to develop custom orchestration around them that bring about compliancy chains and consistent security policies.

When working with other cloud platform technologies and hyperscalers always read the vendor's official documentation and follow their guidance and practices when working with their platforms.

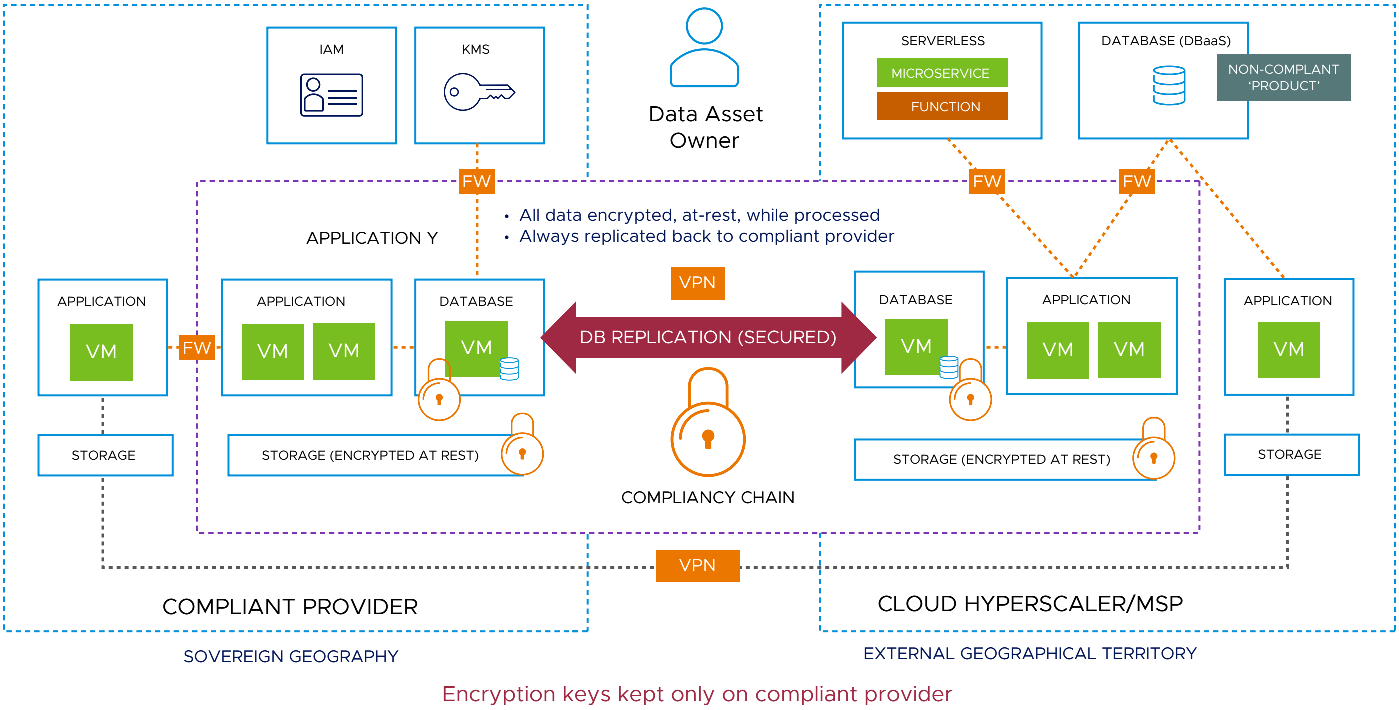

Figure 15 - The Sovereign Cloud Compliancy Chain

Conclusion & Next Steps

This white paper has discussed some of the considerations and concepts involved in developing a secure, trusted Sovereign Cloud platform that can protect the sovereignty of data. Sovereign Cloud platforms are primarily intended for use by cloud and hosting providers serving Government, Military, and private sector customers but can also be implemented directly.

For more information on how to design and build a Sovereign Cloud or to find out more how VMware can help, please contact your VMware Account Manager.

Appendix

Acknowledgements

It is with gratitude that the author acknowledges the following people who invested their considerable time and expertise in helping to get this document prepared.

Juergen-Markus Sobotzik – Enterprise SPP/EMEA Program Manager, VMware

Alex Tanner – Senior Staff Cloud Solution Architect, VMware Cloud Provider Program

Michael Crowley – EMEA Industry Director, Public Sector and Healthcare, VMware

Martin Hosken – Chief Technologist, Cloud Services, VMware

Cory Allen – Staff Technical Product Manager, VCPP Engineering, VMware