Achieving Security Integrity in Service Provider NFV Environments

WHITEPAPER

VERSION 1.1

A Growing Awareness of Security Risks and Threats

The recent increase in sophisticated, targeted security threats by both insiders and external attackers has increased the awareness and urgency from communication service providers for comprehensive security strategies. Service provider networks are particularly vulnerable given the vital role they provide in interconnecting all aspects of society.

Inadvertent actions or malicious abuse by an insider is the most common source of security breaches. Insiders are generally trusted users such as employees, contractors and business partners. Often, well-intentioned insiders may be responsible for costly service disruptions (as a result of inadvertent network configuration changes) in addition to the source of security vulnerabilities. Malicious insiders are often disgruntled employees or contractors driven by financial or personal motives and often seek valuable or sensitive information that can be used to harm the organization.

Advanced threats from external attackers are increasingly being led by organized groups such as professional criminals attempting to gain access to valuable or sensitive information, state-sponsored groups engaged in industrial espionage or cyber warfare, and "hacktivists" pursuing a variety of social causes. These groups are all highly motivated, technically advanced, and are often well funded.

FIGURE 1: MOST COMMON CAUSES OF SECURITY BREACHES

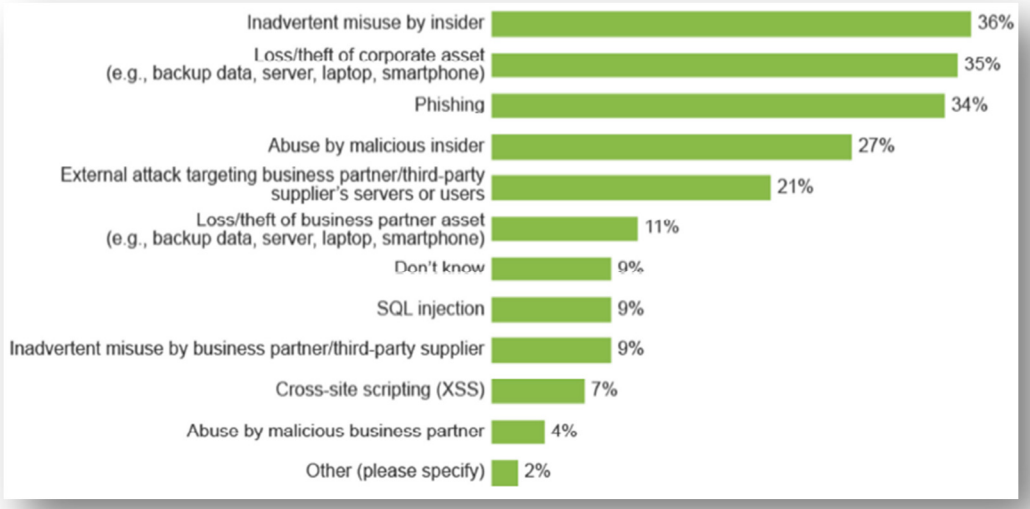

Source: Forrester Research, 2013

Significant Commercial Consequences

According to a recent IBM study, the average cost of an IT/telecommunications outage is $53,000 per minute of downtime. Further losses due to reputation-related costs can add up to millions of dollars. Analysis by the Ponemon Institute reveals that the average total cost per data breach and cyber-attack is $3.5M and this number is increasing annually. Companies report costs ranging from $1M to $60M to resolve these incidents. Severe cases from criminal intrusion or misuse can result in staggering costs into the billions of dollars, not including additional economic impact result from

damaged reputation. Costs and consequences vary by industry; with communication service providers particularly impacted.

FIGURE 2: COMMON THREATS IN TERMS OF ECONOMIC IMPACT

Source: IBM Global Technology Services

Governments and industry groups continue to enact new legislation and compliance standards regarding data protection and privacy and internal control. Compliance requirements continue to become more stringent in response to the complex and evolving threat landscape. Some examples include, but are not limited to:

- Payment Card Industry Data Security Standard (PCI-DSS)
- Federal Information Security Management Act (FISMA)
- National Institute for Standards and Testing (NIST) Network Security Standards
- North American Electric Reliability Corporation Critical Infrastructure Protection Plan (NERC-CIP)
- German Federal Financial Supervisory Authority (BaFin) Minimum Requirements for Risk Management
- Indian Department of Telecommunications, Telecom Enforcement, Resource and Monitoring (TERM) Cells

The common theme of these regulations is generally a requirement that organizations implement security controls over user accounts and access, establish accountability to specific users, and maintain a recording of certain sessions.

Organizations experiencing substantial data breaches might even be subject to fines from regulators if found guilty of negligence. An Institute for Risk Management survey reveals that fines may range from <£50,000 to >£250,000 for British companies who have their data breached to the detriment of the public. In different countries, significant incidents may also be subject to fines by regulators, with penalties reaching into the millions of dollars.

Service Provider Challenges

The business challenges faced by communication service providers are well known: revenues remain threatened by increasing competition from both traditional, and new, more agile competitors. There are urgent needs to develop new service offerings to drive new revenue growth, while improving operational efficiency shrink costs. As a result, service providers are turning to new technologies such as Network Function Virtualization (NFV) and Software Defined Networking (SDN) along with new business models, as mechanisms to increase revenue while reducing operating costs.

In the past, networks were built on customized and purpose built hardware and software. The technical knowledge required to compromise a given network was contained within a very small community, and there were limited motivating factors to encourage hacking. Similarly, network equipment and associated functionality was vendor- and domain-specific, with a small number of entry points, further limiting its exposure to external security risks. Today, the move to all-IP networks has resulted in mobile and wireline networks becoming more vulnerable and exposed to the same types of threats that afflict any server reachable over the Internet.

SDN, like many new technologies, introduces security challenges. SDN involves the functional separation of control and forwarding planes. Securing the interfaces between centralized SDN controllers and the underlying network elements or network functions is crucial to ensure that rogue, malicious instructions dictating how traffic flows across networks are not injected.

NFV completely changes how networks are designed, built and managed. It pulls the functions necessary to run networks off of proprietary hardware and places it on servers that can be deployed where they are needed most – in data centers, mobile base stations, as well as customer premise locations. This combination reduces cost but also dramatically increases the attack surface for security attacks, and increases security administration costs and complexity.

Emerging Security Complexity from New Technologies

NFV is a transformational technology being embraced by service providers. Cost improvements, operational efficiency, and accelerated new service introduction times are some of are some of the market drivers as to why NFV is an integral evolutionary step for most service providers. Some of these efficiencies are achieved by optimizing equipment and infrastructure costs through consolidation of network functions, while exploiting economies of scale from the IT industry.

The standards body ETSI has defined a generic architecture for NFV including management and orchestration (MANO), virtual network functions (VNFs) and virtual network infrastructure (NFVI) as illustrated below.

FIGURE 3: ETSI NFV REFERENCE MODEL

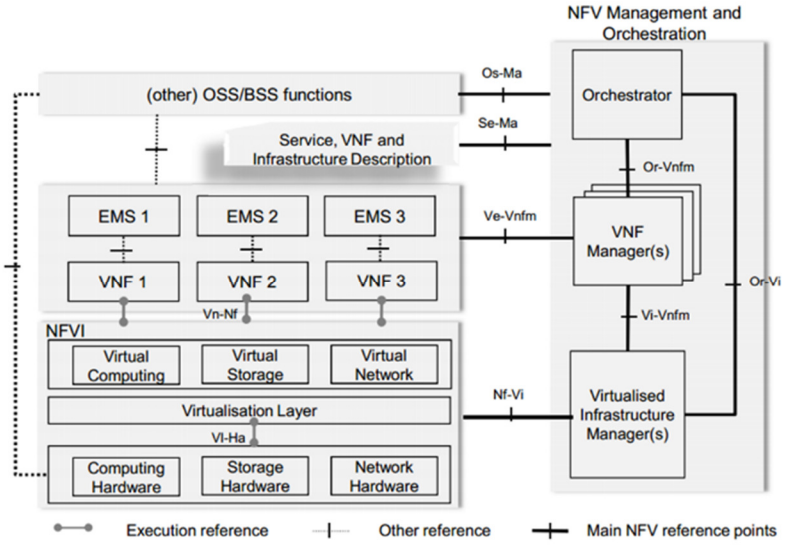

NFV allows service providers to deploy network functions as virtualized software instances instead of dedicated hardware appliances. These software-based network functions can then be driven off of industry-standard high-volume servers, network, and storage platforms. These can be located in data centers, distributed central offices or points of presence, mobile base stations and customer premise locations.

Multiple Administrative Domains

Unlike traditional hardware-based networks, with NFV the hard boundaries that existed between physical network functions are now blurred, making defining and administering security roles and responsibilities more complex. There are multiple levels/domains that need to be addressed, such as the NFVI (including the hypervisor), VNFs, and VNF managers and orchestrators, as well as external systems such as OSS and policy systems. Administration of roles, responsibilities and privilege levels will become more critical and challenging in this complex environment

NFV purists believe that management of virtual networks will be simpler. In some ideal future state, humans will never have to log into networks. The network and communications will be automated, resulting in improved security outcomes. Conversely, pragmatists believe this is unlikely and humans will continue to control processes and resulting changes to networks. The likely reality is that automation will increase but humans will still need to access network resources manually. Configuration errors will continue to occur, provisioning issues will continue, and troubleshooting complex issues will still require human correlation. The definition of identity access management needs to evolve to encompass both people and processes. Administration of who, or which system can view, set, or change configuration parameters and effect network policies becomes vital. This is especially important given the interdependencies between NFVIs and VNFs, and overall service performance and availability. Moreover, as multiple automated software systems access the same shared pool of network resources, assuring that security permissions and policies do not conflict will be crucial. Software enabled provisioning processes can lead to orchestration vulnerabilities including network configuration exploits and malicious configurations.

FIGURE 4: POTENTIAL NFV SECURITY VULNERABILITY POINTS

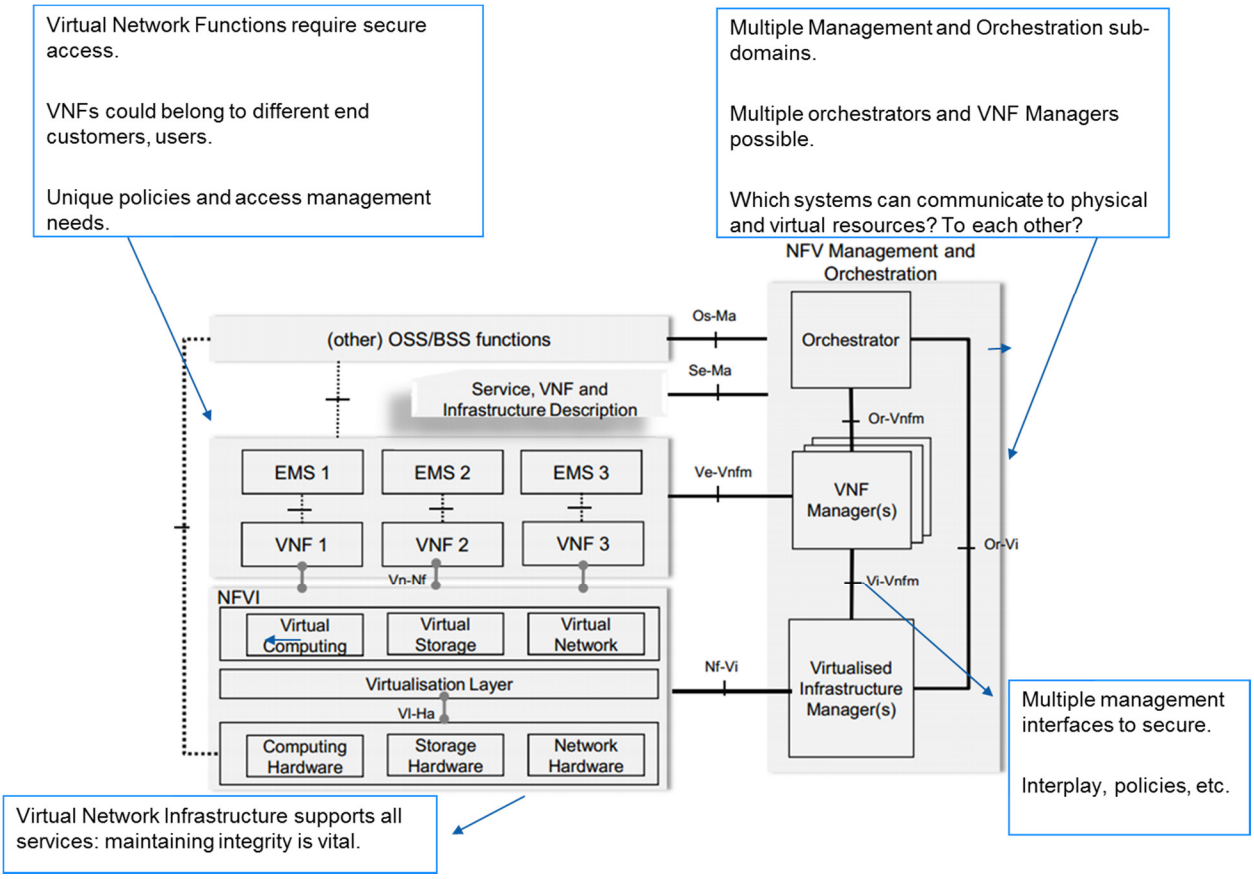

Maintaining Virtualized Infrastructure Configuration Integrity with Multi-tenancy

Multi-tenancy environments pose significant challenges when trying to maintain configuration integrity, and common cloud infrastructure could easily have hypervisor vulnerabilities introduced as a result of integrity failures. Virtual Machine, guest OS, or VNF manipulation could also compromise the integrity of the hypervisor. It will be important that logging and monitoring of hypervisor activities be performed. Similarly, it will be important that VNF configurations themselves are audited to understand whether configuration or operating system changes may have an impact to security integrity.

An important driver for NFV is to create a more flexible and elastic network to enable new service provider business models and revenue opportunities. VNFs will be instantiated, retired, or moved in a more dynamic fashion in order to meet the service delivery requirements. New business models could include VNF, or VNF-as-a-Service, whereby service providers could host different 3rd party VNFs within their own distributed, virtualized infrastructure. Some NFV implementations may involve hosting VNFs from different 3rd parties within a common service provider virtualized infrastructure. Without periodic integrity auditing, VNFs could be arbitrarily instantiated by Virtual Infrastructure Managers on suitable or available hypervisors. This could create vulnerable co-residency scenarios should the hypervisor become exploited or the security policies not be applied properly to the respective VNFs.

Retiring or removing VNFs is equally critical as some VNFs inadvertently left instantiated could result in security breaches or result in susceptibility to Denial of Service attacks. For instance, VNFs may be instantiated for temporary troubleshooting or service testing during service activation. These may include virtual test agents, traffic generators, virtual taps, and packet analysis. If they are mistakenly left instantiated or fail to be retired by an automated process, they can be exploited maliciously or inadvertently during routine network maintenance, resulting in service disruption and extended operational expenses.

Clearly, maintaining configuration integrity will be necessary in order to meet regulatory and compliance requirements, which will be increasingly challenging and potentially expensive in virtualized networks.

Expanded Attack Surface

The aforementioned scenario becomes even more challenging when multi-cloud or multi-site NFV is considered. Technology miniaturization and cost improvements, combined with latency sensitive application requirements, results in a network-wide distribution of virtualized computing and storage assets. The virtual network may span from data centers, remote points-of-presences, to mobile base stations, and to customer premise locations. Not all VNFs are suitable to be centrally hosted for a variety of reasons, including latency, bandwidth and performance. The resulting architecture is very effective and practical for hosting various types of VNFs and changes the convention definition of a security perimeter.

FIGURE 5: DISTRIBUTED CLOUDS EXPAND ATTACK SURFACE

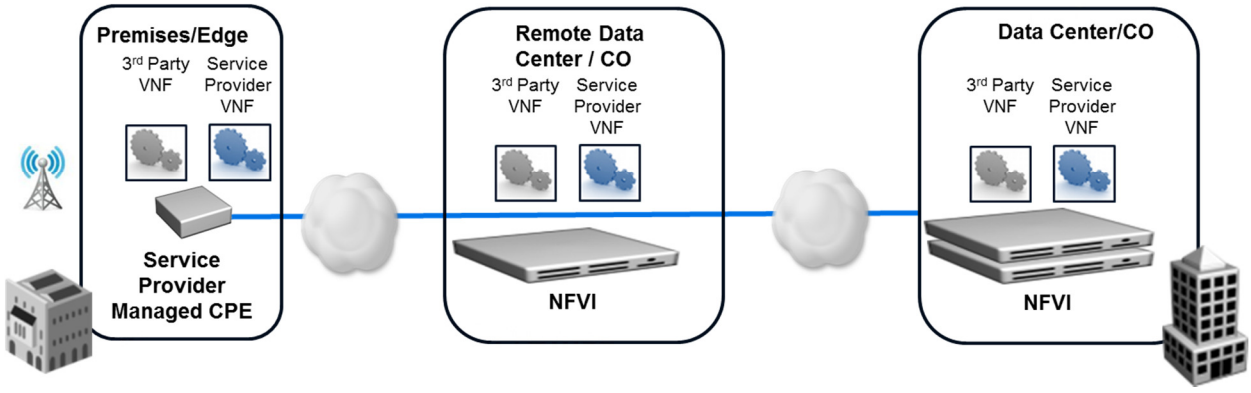

Perimeter-based threat protection approaches such as network, web and endpoint security will be insufficient in these increasingly complex and sophisticated network environments.

Hybrid and Distributed Networks

In the ETSI NFV model, part of the Management and Orchestration (MANO) role is to provide lifecycle and VNF management. Many VNFs themselves will be security virtual appliances. While it is clear that a NFV orchestrator will manage these like any other class of VNF, it is less obvious how security orchestration may be implemented for the VNFs, for NFVI, as well as for the OSS, BSS, EMS and MANO components themselves including orchestrators, VNF managers, and Virtual Infrastructure Managers. For instance, many service providers are envisioning domain-specific orchestrators (e.g. mobile/wireline or business/consumer services may have their own unique orchestration implementations) or a

federation or orchestrators may be necessary simply because of scale (e.g. metro, regional, national and global networks may have unique orchestrators).

Ultimately, end-user services will traverse a combination of networks (e.g. mobile and wireline), regions (e.g. metro, regional, national and global) and technologies (e.g. traditional physical networks and virtual networks). Services and service chains will be complex, spanning shared infrastructures, physical networks, locations, and clouds. It will be important to administer and maintain service-oriented security privileges and policies to ensure that the right systems, processes, and people have the appropriate access end-to-end in order to turn on, manage, optimize, and troubleshoot services.

Maintaining security integrity needs to be part of a holistic service assurance strategy, requiring a service-driven and contextual view of security access control policies.

Organizational and Business Processes Complexity

Service provider networks are inherently complex and heterogeneous:

- Multiple services: consumer, commercial
- Multiple markets: metro, regional, national, global
- Multi-domain: mobile, wireline, content
- Multi-technology: cable, fiber, copper, 3G/4G/5G, Wi-Fi
- Multi-vendor: many unique hardware and software suppliers per domain, per technology
- Multi-generational: legacy technologies and products
- Multi-protocol: IP, MPLS, Ethernet and legacy protocols such as TDD, ATM and Frame Relay

The network complexity is often mirrored by complex operations and business processes, with distinct management and operations silos. Compounding this is the financial pressure to reduce operational costs. As a result, outsourced managed network services are increasing in popularity. Increasingly, service providers are outsourcing network installation, field operations, and network operating centers. In some cases application and service delivery are also outsourced. As a result service provider networks, equipment and systems are being accessed by not just employees but a myriad of third parties, including partners, suppliers, and customers.

Dynamic Access Policy Management

Even granting user access to a wide range of resources through strong authentication measures still poses risks. Once granted access, how long can a user remained logged in? If the user is part of a partner, supplier or other 3rd party it may be more challenging to enforce security best practices-at a minimum, it will increase the cost of creating, maintaining and enforcing SLA's in the future. Maintaining the integrity of the security policies requires role based identity management. Identity access management policies are necessary to control access to resources depending on the type of user and the context of the user access request: is the user trusted? From where is user access originating? What are roles and permissions?

Rapid Provisioning and De-provisioning of Users

Given the fluid and dynamic nature modern software defined networks, the management of user accounts and privileges must also become agile. User and system access cannot be persistent. It will be important to be able to grant access privileges rapidly and even more critical to disable all access when the association ends. This could be an employee resigning, a 3rd party completing maintenance or troubleshooting, or an automated management and orchestration system completing its tasks.

Recommendations

An integrated Identity and Access Management solution that spans physical and virtual networks, and the associated OSS, BSS and management and orchestration systems is needed. Virtualization introduces multi-tenancy with the need to establish flexible role-based security policies for both humans and autonomous systems. There are now multiple layers of interdependencies (e.g. MANO, VNFs, NFVI, physical network functions), which will drive more complex policies. Domain isolation between these different slices will be needed, along with flexibility to create access management rules as needed. Network implementations, services, and operations practices will vary by operator, by service, by region so access management solutions must be adaptable and flexible.

Single sign-on (SSO) is a critical need. SSO facilitates both ease of use and administration simplicity, but also allows for rapid response and control to change and revoke access. SSO supporting role based identity management is crucial.

Service providers must rank scalability and high availability highly as they operate always-on, mission-critical networks. Service provider networks are huge and diverse. Identity access management strategies need to accommodate a variety of network equipment, multiple generations of technologies, and scale to support hundreds of thousands equipment types, virtual infrastructure, virtual network functions, servers, and systems. Conventional systems designed for enterprise applications often lack the scale, performance, and availability needed for service provider networks.

Network behavior analysis will be an integral part of next generation security strategies. Autonomous management and orchestration processes will result in more dynamic and fluid networks. Virtual network infrastructure configuration changes will be frequent, virtual network functions will be instantiated, retired, changed, and moved. The automated systems and humans accessing network resources to activate, change, monitor and troubleshoot services will grow. It will be vital to correlate network configuration and service parameter changes with security events. This will be crucial to pin-point configuration changes which may create security risks and to rapidly identify network access resulting in malicious attacks.

Conclusions

In order to realize the full commercial benefits of new technologies such as NFV, identifying and overcoming some of the practical and critical operational considerations will be required. Security is one critical operational aspect that will quickly come into focus. Securing NFV must not be an afterthought if full benefits are to be realized. The definition of Identity Access Management must evolve extend to systems, as well as people, and take into account the expanded and fluid attack surface.