


SECURITY RESHAPED IN THE DIGITAL TRANSFORMATION ERA



Mohamed Sohail

Senior Engineer
Solutions Architecture
Dell Technologies
Mohamed.sohail@dell.com

Nour Mahmoud

Assistant Professor of Information Technology
Faculty of Computers and Artificial Intelligence
Cairo University
Nourmahmoud@cu.edu.eg

Robert Lincourt

Robert.lincourt@dell.com

Mohamed Hamed

Assistant Professor of Information Technology
Faculty of Computers and Artificial Intelligence
Cairo University
Mnasrtaha@cu.edu.eg



The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged/Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage
- Enterprise Architect

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

Learn more at www.dell.com/certification

Table of Contents

Digital transformation and new attack surfaces	6
Democratized access to infrastructure and data	7
Security becoming more opaque	7
Threat modeling	8
The Data Decade is upon us	8
Why we should worry about Digital transformation	9
Phases of a cyber-attack.....	9
Shortening your defending OODA loop having a blind spot	10
Collection isn't detection.....	10
A) We need to reduce manual steps and manual errors. The more steps we have to physically take, click on, look at, and search for, the slower we can react to an attack.....	10
B) Implementation of investigation tools can help eliminate slow detection of an attack.	10
Attack surfaces expansion.....	11
Most common and important attack surfaces.....	11
Device security	12
Measurements in the project.....	13
Challenges ahead.....	14
The solution.....	15
Novel approach #1 - Mechanism of anomaly detection	15
Novel approach #2 - Employing the trust scoring methodology.....	16
Establishing power behavioral profile.....	17
Benefits of this approach	17
Platform security	18
Multi-cloud security model from Defensive to Autonomic.....	18
What is Multi-cloud Environment?	18
Benefits of Multi-cloud.....	19
Disaster Recovery	19
Accommodating Peak Usage	19
Avoiding Vendor Lock-in.....	19
Allowing You to Pick and Choose	19
Security Perspective from Defensive to Autonomic	19
Dell.com/certification	3

Defensive Security	19
Key Capabilities.....	20
Distributed Security.....	20
Efficient Security.....	20
Key Capabilities.....	20
Agile Security	21
Key Capabilities.....	21
Autonomic Security	21
Multi-cloud Security	22
Challenges of Multi-cloud security	22
Best Practices for Multi-Cloud Security.....	23
Multi-cloud security use case “MUSA”	24
Main framework components.....	25
Security fragmentation.....	27
Conclusion	29
References.....	30

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell Technologies’ views, processes or methodologies.

“My money is in limbo” – said Travelex customer Travelex: Banks halt currency service after cyber-attack

Can we really imagine seeing one of the well-known money services companies in 2020 forced to use pen and paper due to a cyber-attack???!!! Yes, we can....This happened with Travelex at some banks and supermarkets. Travelex had said little publicly since hackers held its systems to ransom by encrypting its digital files, reportedly demanding \$6m (£4.6m) to unlock that data.

This turned out to be one of the largest computer system outages as ransomware attacked the core systems of the main operations, leaving thousands of customers stranded, with customers and banks that deal with the affected agency waiting for clearance.



However, not paying can be extremely costly. Steel producer Norsk Hydro was hit by the LockerGoga ransomware last March. Some 170 factories and offices were taken offline, with manufacturing partially suspended. The hackers demanded an estimated £300,000 but the company instead refused to negotiate and has spent about £50m recovering operations.

“The crux anxiety for most of the modern Data center managers is: hardening devices against intrusion is a good first step, but it is nowhere near a complete and holistic security model specially for the backup’s healthiness.” As a result of the intelligence of the modern cyber-attacks that tackle the backup data, there is a stringent need to explore new directions to protect and build a solid strategy to leverage a secure data stream network and its accompanying services to provide enterprise-level end-to-end security. Doing so shifts the primary burden of securing billions of new devices from hardware manufacturers into the network layer, which is far more flexible and robust for ongoing security, plus adding a new way to build a strategy to lessen the damage and operate seamlessly in case of a disaster.

According to reports from tech industry leaders such as Cisco and Microsoft, cyber-attacks against businesses of all sizes are on the rise and doing more damage than ever. No one can ever ignore that businesses can be at risk of losing important data. Lost data leads to costly downtime, customer dissatisfaction, regulatory fines, and lost revenue. As a result, IT pros must meet extremely high expectations. You need to keep the company running 24-hours a day.

Can we imagine all this just started with vulnerable bug into one of the used operating systems kernel service? To add to the woes, the IT Ops Command Center took more than one week to restore the service and another 3 weeks to check how to deal with the lost data. Apparently, no one was prepared to deal with such chaos or even to predict how fierce this disaster.

By 2021, worldwide cybercrime damage is expected to reach \$6 trillion —double what it cost businesses in 2015. As digital transformation sweeps the globe, the imminent threat of cybercrime grows alongside it. As a result, new techniques in cybersecurity are being developed to mitigate the increased levels of risks.



Digital transformation and new attack surfaces

Sophisticated cyberattacks employ a variety of effective tools and tactics: phishing scams, malware and spyware attacks, browser and software exploits, access through lost and stolen devices, and social engineering. It takes continuous vigilance to maintain visibility across the threats you know, and to become aware of emerging vulnerabilities. While there are certainly tools to help maintain an always-on approach to security, the reality is that security demands a multifaceted approach to ensure your organization is prepared to handle new attacks no matter when they occur or where they come from. To design a comprehensive defense strategy, it is important to understand the methods and tactics employed by modern attackers.



Figure 1: Digital transformation

To be future-ready and support development and deployment of innovative models or applications in the era of digital transformations, organizations need some sort of **Autonomous Operations** (AO) solutions. In the coming decade, applications powered by AI technologies and Intelligent Automation will transform workplaces, people, and interactions between technology producers and end users. This would also require higher cognitive skills such as creativity, critical thinking, decision making and complex information processing across the entire eco-system.

In this Knowledge Sharing article we highlight security best practices that don't slow your business down. Rather, they speed it up along with new inventions in the security domain. We will also show, while being optimistic with the revolution that will come with the digital transformation, why we should worry about security.

Democratized access to infrastructure and data



Figure 2: Data access Democratization

A key driver to enabling digital transformation is providing access to infrastructure and data in new ways. The concept of cloud has proved that a simplified view and access to infrastructure is not only possible but a desirable pattern. Allowing all people to interact with infrastructure and services at a level they are comfortable with provides everyone the ability to innovate at speeds never seen before. This leveling of the innovation playing field is powerful but fraught with the potential to create enormous security attack surfaces.

Security becoming more opaque

Modern datacenters have made huge strides in providing a more robust operating environment for traditional and cloud native apps. The ability to execute an app without fully understanding where all the pieces of the app will execute is a key attribute of multi-cloud.

Creating policies for everything within the datacenter or a multi-cloud environment is complex to say the least. Getting this correct is becoming increasingly difficult. Vendors are now supplying multiple tools to help customers navigate these complex environments. AI/ML algorithms here help identify where policies start to overlap or undo existing checks.

This modern datacenter's topology is changing so fast that images and policies that were once very strong and secure have been replaced by more options or secure features. Ownership of each layer of security is becoming harder to identify and its consistency fluctuates often. Consider a multi-cloud application that is running some services on-premises and some in a public cloud. Each own the physical security of the resources and securing the primary access points and basic security of the images being deployed. Moving up a layer in to the virtualization, connectivity, and application layer, security of these layers start to distribute across the provider (both public and on-premises) and the application development and security teams. Creating a



threat model identifying each layer becomes difficult because looking into the black boxes of the cloud and modern datacenter we only see what is exposed to us.

Threat modeling

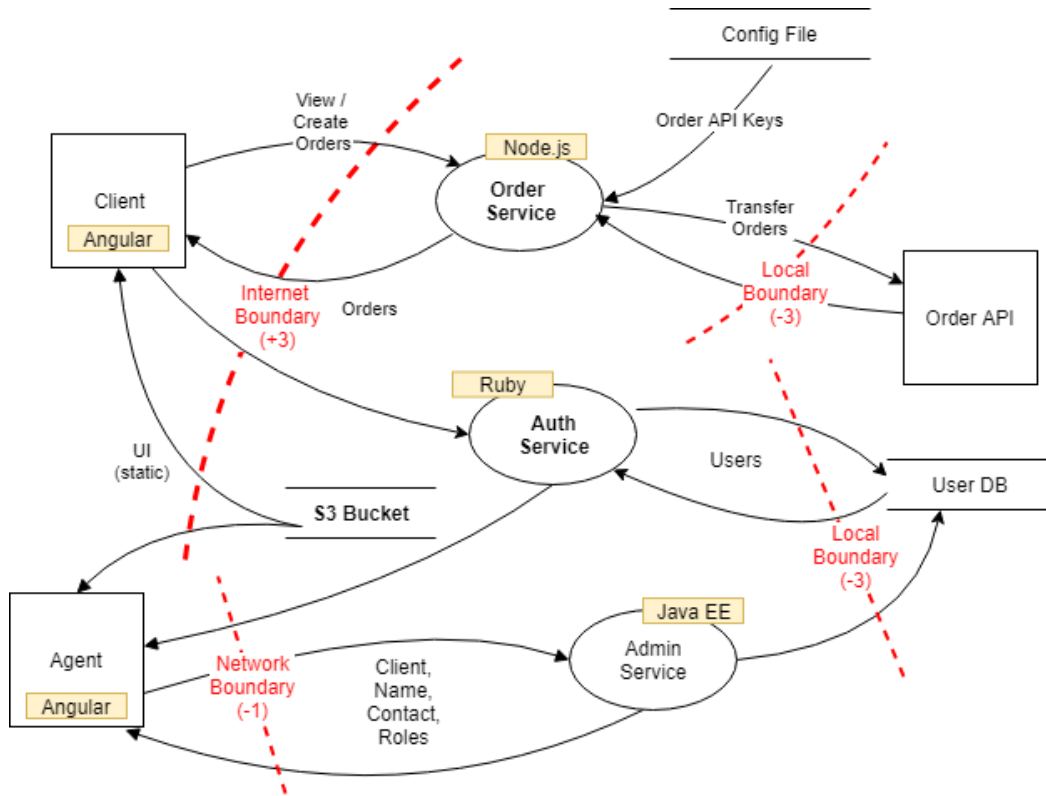


Figure 3: Example of Threat modeling

The threat model is a great tool to identify risk. When this threat model is applied in a more static environment, we can see the mapping of resources to application layers. In this new dynamic environment, mapping is able to constantly change. If the application's front-end web server is now being hosted in a public cloud does the security posture need to change? We have seen examples in the Capital One data breach. If the misconfigured service was hosted on-premises the data would have been exposed to internal employees only. One could argue this might have been identified within a threat model. This exercise would have only helped if all the attributes were visible to all involved.

The Data Decade is upon us

With the promise that AI/ML shows in multiple areas, the rationale for collecting more and more data is understandable. We are starting to collect data in every aspect of our lives and trying to form a quantitative perspective from it.

Remember the issues from democratizing access to infrastructure? Some of the same issues apply to all the data we are generating. Instead of providing access to individual pieces of data, a higher-level policy mechanism will need to be put in place. Again, the policy management issues come into play, so an AI/ML approach of monitoring and granting access to individual data fields is

needed. While you might have access to certain types of data, not all fields might be visible for you.

A primary reason for democratizing access to data is our desire to gain foresight and understanding from it. Data from a single source is interesting and provides a nice structure to identify solutions or create incremental improvements to our ecosystem. However, the ability to corollate data from multiple sources allows for use to alter the ecosystem and provide entirely new experiences.

Utilizing all this data we will have to provide more access to multiple data sources, but at what risk to privacy?!! AI/ML models can be used to validate how data is being used and for what purposes. This could provide more control to the people or devices creating this data.

While most AI models require a massive amount of data, new techniques are being created that needs only a fraction of the massive data set, but the quality of the data required is much higher.

Cleaning and wrangling data becomes the bigger problem in this new decade. Biased data is also a huge problem – more and more examples of this are seen every day. Understanding and publishing these new attributes of a data set is key to enable better experiences with AI.



Why we should worry about Digital transformation

Phases of a cyber-attack

Despite the versatility and velocity of the modern attacks, the majority of them share the same characteristics starting like building a novel. The beginning is “the multi-phased approach”, with the objective to establish the known term “persistence in the target environment” by using the following sequence of steps:

- **Reconnaissance**, during which attackers gather information about the target environment, including its resources and identities. This information forms the basis for subsequent phases of the attack.
- **Lateral movement**, during which attackers extend the scope of their attack inside the target environment.
- **Domain dominance (persistence)**, during which the attackers capture the information allowing them to escalate their exploit using various sets of entry points, credentials, and techniques.

Another pattern that represents how more sophisticated attacks are carried out is OODA, representing four stages in a typical attack – **Observe, Orient, Decide, and Act.**



Shortening your defending OODA loop having a blind spot

Observe, Orient, Decide, and Act.....a main point that needs to be addressed is eliminating blind spots. A blind spot in newly digital transformed organizations, increases the risk of missing an important signal that an adversary can operate in. A blind spot that we don't know is a great threat. For instance, in a SaaS app, you never know what they're doing on your hosts or your network. You've got to have as much visibility as you can fit in.

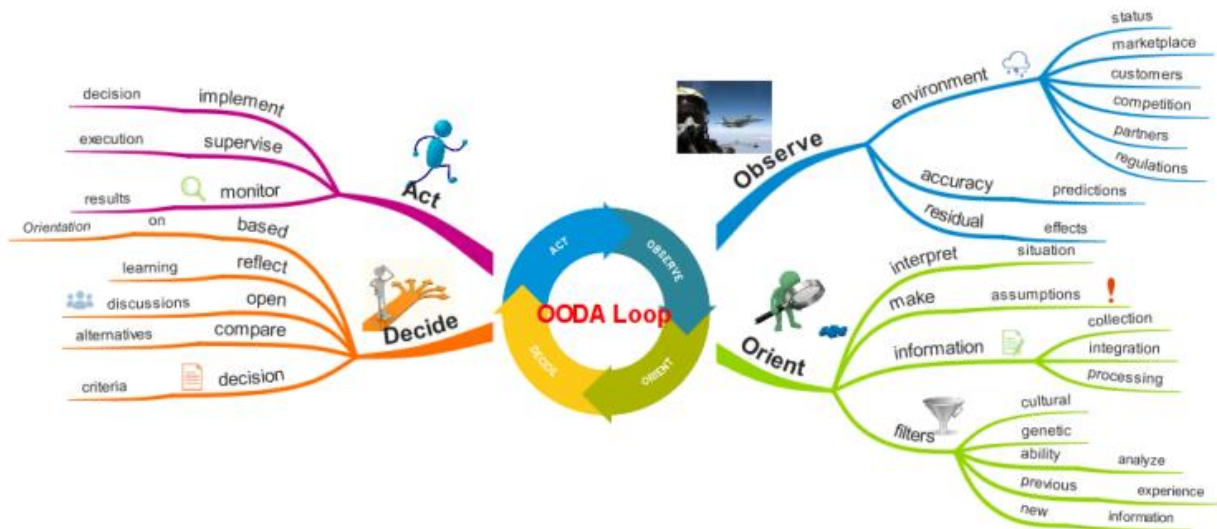


Figure 4: OODA Loop to Accelerate Your Decision Making

Collection isn't detection

Many organizations have great network visibility, but don't have a lot of visibility to what's happening with their identities and their credentials. That's a critical place where adversaries operate in the defender blind spots today. As a next strategy, this principle instantiates in a couple of ways:

- A) We need to reduce manual steps and manual errors. The more steps we have to physically take, click on, look at, and search for, the slower we can react to an attack.
- B) Implementation of investigation tools can help eliminate slow detection of an attack.



Attack surfaces expansion

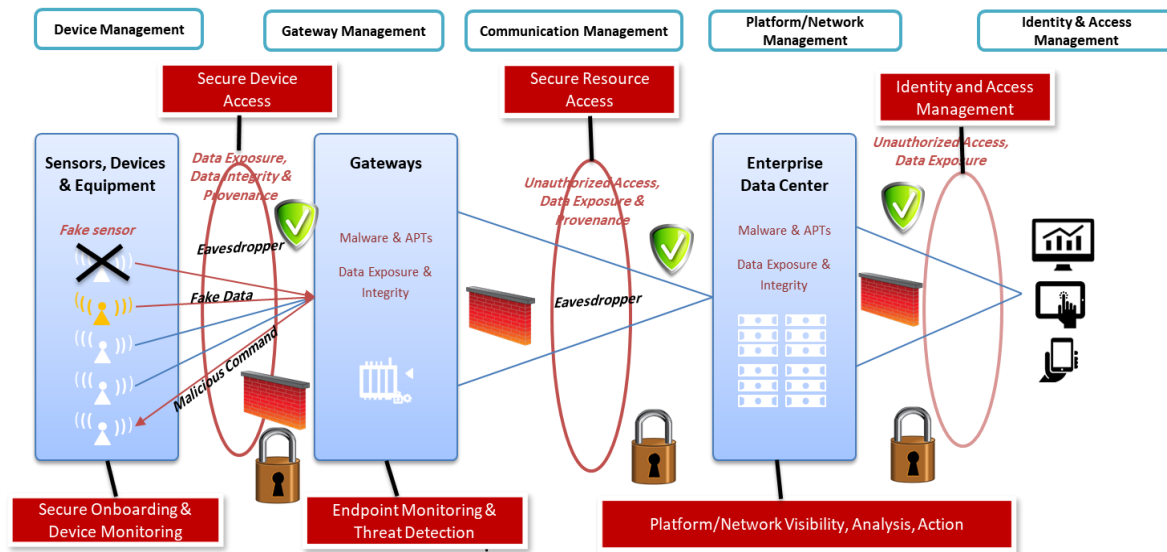


Figure 5: Modern DC Potential attack areas

Two external factors are forcing organizations to rethink how they protect and store their data, especially in the digital transformation era. The first is the ever-increasing threat of ransomware, which encrypts production data forcing the customer to pay for a “key” to decrypt it. Rapid cyberattacks like Petya and WannaCrypt have dramatically changed expectations regarding the speed and scope of resulting damages. Rapid attacks have extremely high propagation speed. This leaves little time for defenders to react (detect + manually respond or detect + write automatic response rules), underscoring the importance of preventive controls and recovery processes. In 2017, among the global enterprise customers, these rapid cyberattacks took down most or all IT systems in about one hour, resulting in \$200M – 300M USD of damage at several customers. The second is data protection and privacy regulations, such as the European Union’s (EU) General Data Protection Regulation (GDPR).

Most common and important attack surfaces

The following illustrates in depth the most common and important attack surfaces that are a golden target of attackers, and how to redesign your data protection strategy either on-premises or in the cloud for governance and compliance. The 3 main topics are:

1. **Device security:** where the data is generated.
2. **Platform security:** Where the data is exchanged and stored.
3. **Application / monitoring security:** where the data is presented.

We will discuss new strategies of the new datacenter regulations, and how they affect data growth and Cloud adoption as well as new interesting novel ideas that address unexpected problems and the solution for it




Figure 6: Attack targets

Device security

Leveraging a new security paradigm for tackling security breaches through Power consumption trust score.

“The crux concept for IoT manufacturers is this: hardening devices against intrusion is a good first step, but it is nowhere near a complete security model.” The strategy that we propose in this knowledge sharing article is to leverage a secure IoT ecosystem based on a novel approach by leveraging the potential of studying the sensors’ operational behavior to provide enterprise-level end-to-end security for IoT devices. Doing so shifts the primary burden of securing billions of new devices from hardware manufacturers into a security layer, which is far more flexible and robust for ongoing security.

We propose a **distributed IoT security paradigm** using trust scoring and authentication leveling applying data analytics techniques, trend analysis and anomaly detection algorithms. Our proposal aims to implement a simple and easy approach to achieve higher security by leveraging the IoT devices themselves (in a crowdsourcing fashion) to provide trust scoring and authentication leveling making it easy for legitimate users but hard for intruders.

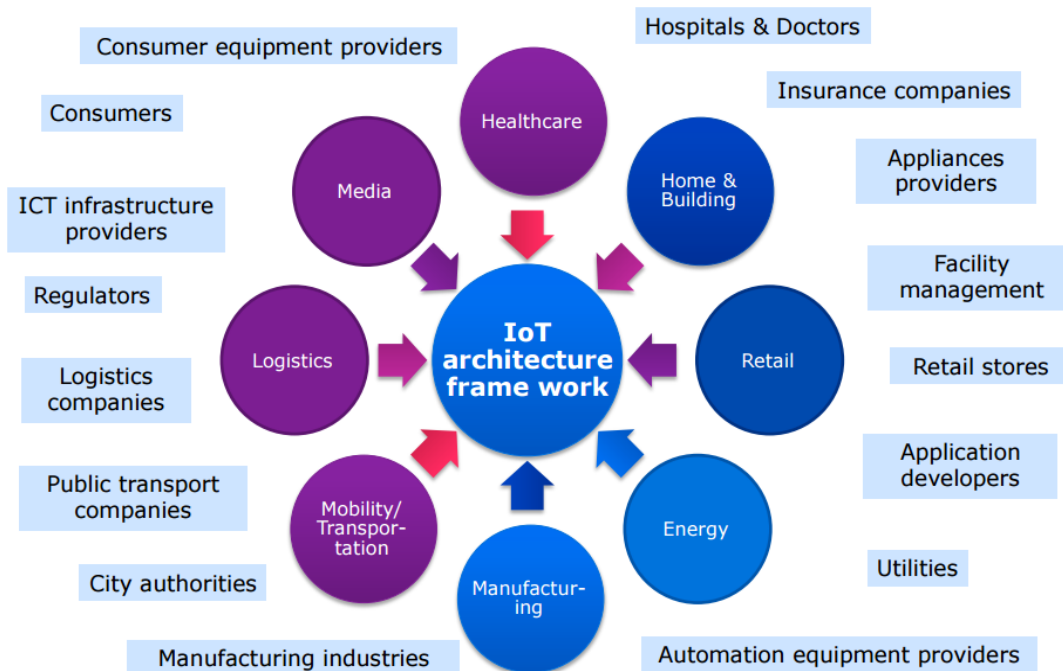


Figure 7 : IoT Application Domains & Stakeholders

Suspicious behavior of sensors and hyper activity are major meters that we can rely on to detect potential threats within a distributed environment. The problem we solve is the proper detection and blocking of suspicious and abnormal activity. By adding certain algorithms to the chips of the sensors, we can detect and allow legitimate activity of sensors and block others, based on a new mechanism of power consumption trust scoring.

Measurements in the project

In this approach it is crucial to be pragmatic on helping, designing, and implementing an easy to implement solution, based on an innovative, powerful, sustainable, and world-class level implementation. This comes from investing in an infrastructure-related initiative delivering on a strategic long-term vision. With this vision in mind, we implemented the idea to satisfy the need and eagerness towards a robust security strategy.

In our design, we considered these business challenges:

- Cost Competitiveness
- Highest Levels of Reliability and Availability
- Ease of Management
- High Performance
- Compatibility

Cost Competitiveness

Being cost competitive is paramount in order to build and maintain business. Whether we are facing economic turmoil or economic boom times, we must guarantee a security solution that doesn't incur economic burden and fits the customer's budget.

Highest Levels of Reliability and Availability

Offering cost-effective solutions is meaningless if the solution is plagued by outages, or not performing well. Security implementations must be capable of performing even after suffering multiple attacks and rely on achieving high level of reliability and use the crowd sourcing mechanism. Customers loyalty will be lost if availability obligations are not met.



The more complex the solution, the more resources it takes to maintain and operate over its lifecycle driving overall cost up while driving reliability down. Ensuring staffing levels remain stable in the face of unabated growth is essential in cost containment and is the main reason ease of management remains a key requirement.

High Performance

It is critical to deliver the highest levels of performance even during peak usage periods. Because adding a new security layer generally causes some slowness in the overall solution, it must be capable of delivering during periods of high usage and must be designed to eliminate congestion points. Delivering solutions that suffer from poor performance frustrates customers and wastes precious time and resources tracking down and resolving performance-related issues.

Compatibility

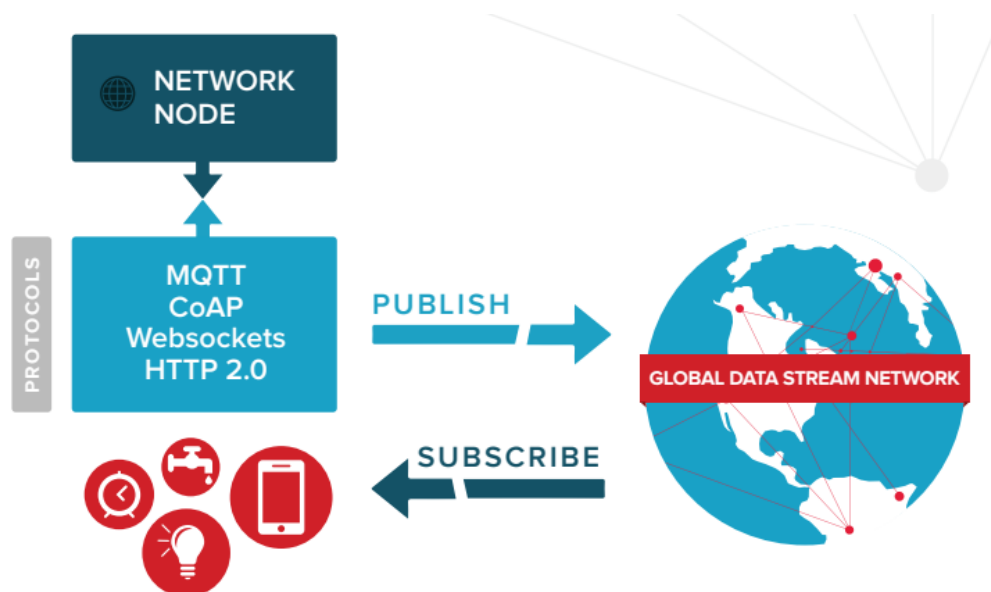
Gone are the days of implementing independent computing silos, especially in the world of IoT, our focus for this use case. To maintain aggressive growth objectives, we need to ensure all of the systems being deployed are compatible and work with one another. It's expensive and difficult to maintain solutions designed in isolation. Everything needs to work together and scale in order to keep the overall solution as simple and manageable as possible. Our proposed solution has been tested in a real-world data center. The architecture has been designed to scale to add other types of sensors and actuators to the deployed environment.



Challenges ahead

“Any device on the Internet with an open inbound port will be attacked. It's a matter of when, not if.”

For one device – i.e. a server – to push data, another device (i.e. an IoT device) has to be listening. In a traditional model, the listening device will open an inbound port and wait for data to be pushed. While this can work in some scenarios, it is a massive risk for IoT as these ports must remain open indefinitely. The security risks of leaving inbound ports open include malware infections, modification or theft of data, DoS attacks, and arbitrary code execution.



Device's nature

IoT devices are often released and NEVER upgraded, which means they become more and more vulnerable to attack as time goes on. Additional functionality and features such as security, unfortunately, are not included within the devices themselves. We need a way to detect such attacks.

Upgrades and patching limitations

Common IoT devices, such as thermostats, garage door openers and even alarm systems, are typically small form factor devices with very little surface area where chips or other devices can be installed. As a result, only basic functionality such as reporting, monitoring and alerting are included within their programming.

Behavioral profiling

We should have the vendor create a “power profile” with expected power ranges for different device activities. This can become the baseline against which the analytics consult.

Management standardization

Security techniques rely heavily on the devices themselves which adds complexity and dependency on the manufacturer hardware and software. Security needs to be decoupled as a separate component that can be standardized and seamlessly managed.

Current security techniques are not dynamic. Rather, they concentrate on solving the issue after it happens, not taking into consideration the valuable information that can be obtained from the devices.

The solution

Consider this demonstration of a detailed breakdown of the various areas. Our contribution objective is based on the already approved patent US10097572B1, to leverage the power consumption as a trust scoring mechanism to detect threats on the IoT ecosystem.

Novel approach #1 - Mechanism of anomaly detection

Power consumption and its relation to security paradigms can be an important step towards tightening IoT security measurements. Figure 8 illustrates how the mechanism goes to detect any anomaly within the network based on the amount of consumed power of the sensors’ nodes. In this algorithm we propose having the vendors create a “power profile” with expected power ranges for the different device activities. This will be the baseline to build the behavior profile for every sensor’s node based on the reported power usage by every node in different scenarios.

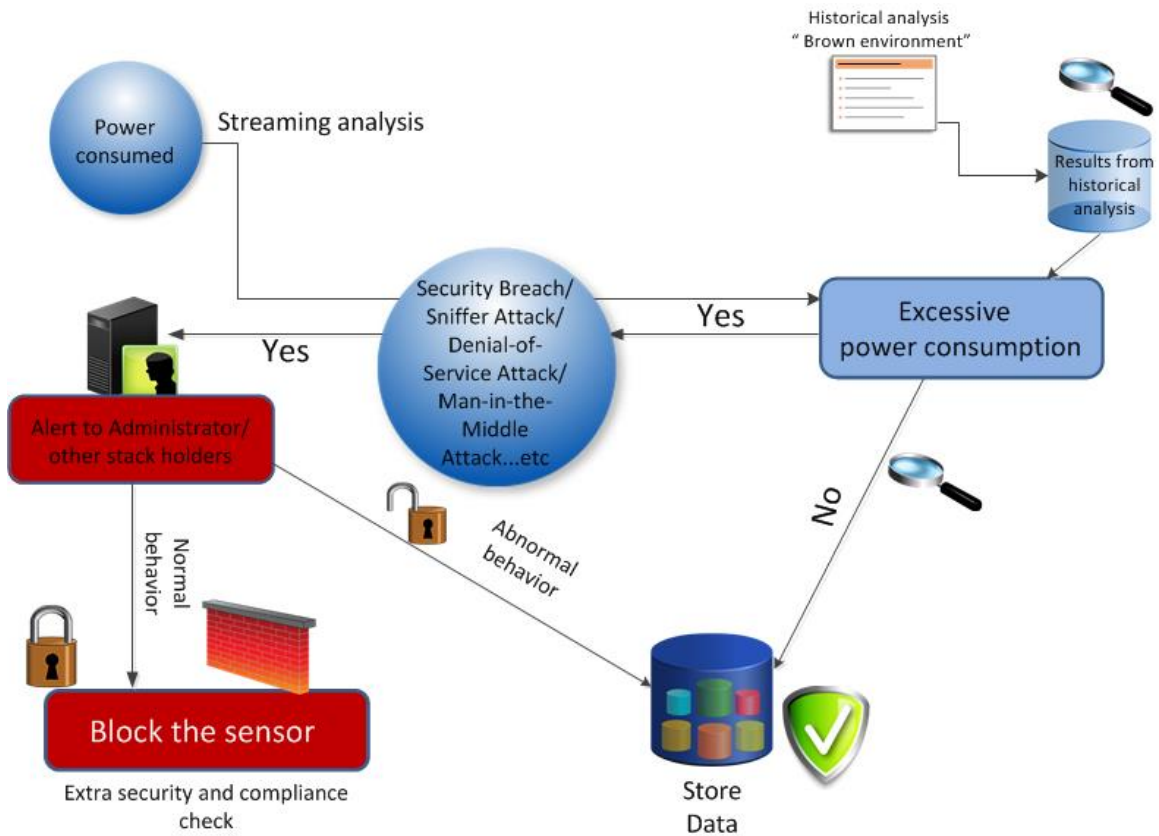


Figure 8: Sensors' behavior algorithm

Our proposed ideal learns the holistic behavior of this interplay between the sensor nodes in a way that gives the administrator a proactive hint to detect an attack.

Novel approach #2 - Employing the trust scoring methodology

The approach will use the same intelligence that enables devices to perform their tasks but must also enable them to recognize and counteract threats. Fortunately, this does not require a revolutionary approach, but rather an evolution of measures that have proven successful in IT networks, adapted to the challenges of IoT and to the constraints of connected devices.

We define the architecture and the underlying algorithms to be used as well as the market use cases and solution packaging.

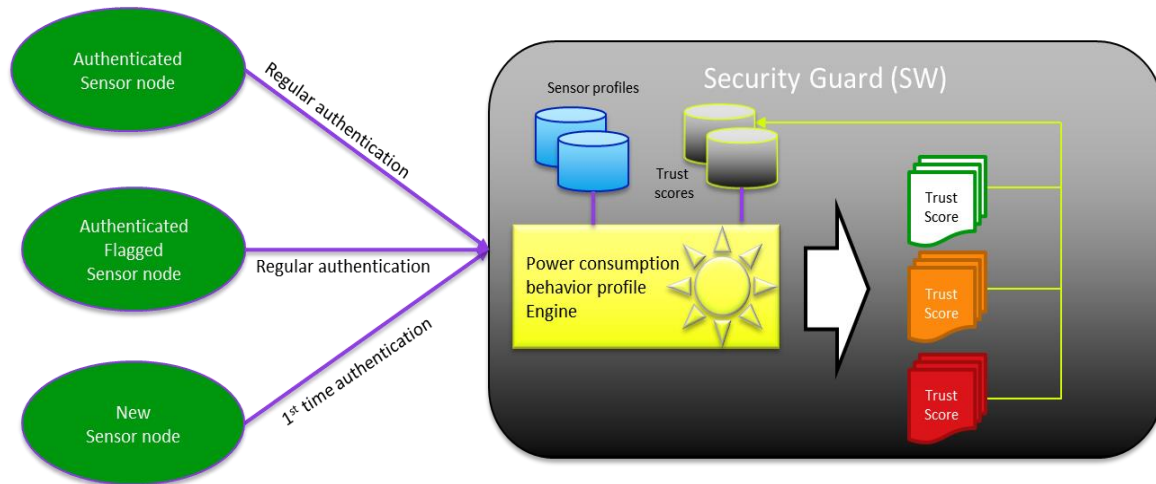


Figure 9: Power consumption trust scoring mechanism

Establishing power behavioral profile

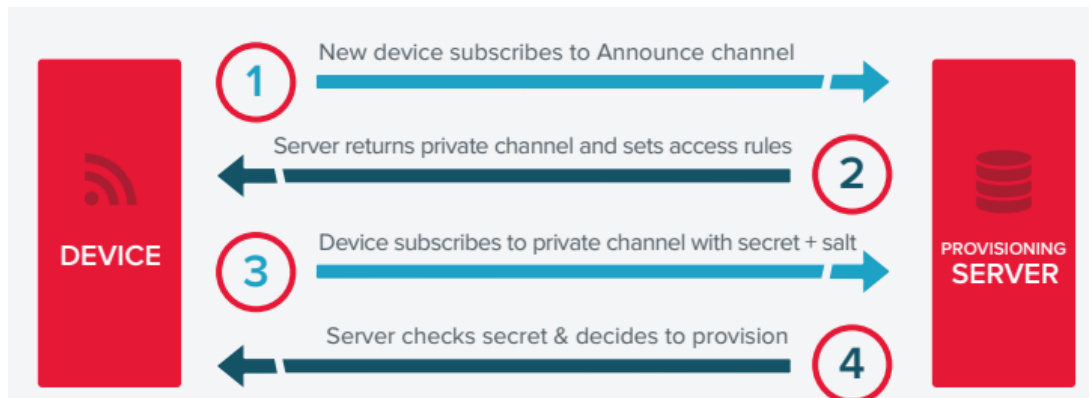


Figure 10 : Securely provisioning IoT devices

Power Behavioral Profiling, and Anomaly Detection for devices with ACPI integration

The power behavioral profiling mechanism follows the interactions of human operators with the system and learns the normal operator procedures for the used power of every sensor's node. This mechanism will issue an alarm when deviations from normal operator behavior occur. It also detects behavior patterns unique to individual operators, and sensors too.

Benefits of this approach

- No need to change the current infrastructure of the solution.
- Can be embedded as add-on technique to harden the current security solution without loading the system with extra software.
- Adds a proactive security layer for the current solutions.

- Next generation Data Centers (DC) will be fully connected and serve mostly IoT-related activities. Power consumption information needs to be included in the DC vulnerability assessment and risk mitigation.
- Implementation of this solution can be integrated into edge management tools such as VMware Enterprise mobility management product suite.
- This approach is designed for the ‘higher orders’ of IoT where processors can be instrumented to gather the necessary telemetry data.
- Leverage IOT Gateways such as Dell gateways for deployments.
- With the use of Machine Learning and Deep Learning, this solution may also be implemented at the core/DC level or as a service in the Cloud. Deep learning enables continuous learning with the deployment of retrained neural network.



Platform security

Multi-cloud security model from Defensive to Autonomic

Organizations have an opportunity to evolve their capabilities to adapt to changes in IT infrastructures and a move toward *multi-cloud environments*. With this security evolution, they can not only maintain their security posture in the face of rapidly changing technologies but also shift the odds in their favor against attackers. Several factors arise in this evolution, including leveraging “home court advantage” for deception techniques, sharing of inter- and intra-industry threat intelligence, use of automation and autonomies to optimize the efficacy of the organization’s security personnel, and others. What follows is an overview of the pathway toward *multi-cloud security architecture*, rooted in the evolution of organizations’ security.

What is Multi-cloud Environment?

Multi-cloud is the use of multiple cloud computing and storage services in a single heterogeneous architecture. This also refers to the distribution of cloud assets, software, applications, etc. across several cloud-hosting environments. With a typical multi-cloud architecture utilizing two or more public clouds as well as multiple private clouds, a multi-cloud environment aims to eliminate reliance on any single cloud provider.



Organizations increasingly want to take advantage of the flexibility and choice of multiple cloud offerings to use best cloud services while achieving satisfactory cost reduction benefits. International Data Corporation (IDC) predicted in their last annual report that multi-cloud adoption will increase drastically, citing that more than 85% of enterprise IT corporations will invest in multi-cloud architecture by 2020.

With the enterprise hurtling towards digital transformation at breakneck speeds, adoption of cloud – both public and private – has accelerated. Plus, working across different clouds across various platforms brings a wide range of challenges in its wake. Lack of understanding of cloud technology is the most basic one which CTO's face.

Benefits of Multi-cloud

Disaster Recovery

It gets risky when an organization uses one cloud platform to manage all organizational resources. A cyber-attack can take down all the operations for a long-time leaving end-user inaccessible until it resolves. Using multi-cloud architecture makes your company's services resilient against these types of cyber-attacks because there are other clouds available to take on the workloads when one cloud goes down.



Accommodating Peak Usage

Another reason an organization may choose to go the multi-cloud route is to enable cloud bursting. This means that applications on one cloud platform can burst temporarily to another already-in-place cloud platform when the need for computing capacity peaks and for economic efficiencies.

Avoiding Vendor Lock-in

Relying on a single vendor for their products and/or services and the inability to move to another platform without incurring hefty fees is notoriously known as “vendor lock-in.” Lock-in may be directly enforced by the cloud provider or could be due to technical issues and dependencies.



Allowing You to Pick and Choose

Using a multi-cloud approach allows you to pick and choose the best of what each platform has to offer. It enables you to create a customized, flexible solution to meet your needs. For example, your organization may wish to use certain machine learning developer tools offered by AWS but prefer Google's high-speed database services. Multi-cloud gives you the freedom to pick the best components from each provider to create your ideal setup.

Security Perspective from Defensive to Autonomic

Defensive Security

In a Defensive Security level, the primary focus is to keep attackers out using perimeter defenses designed to stop externally-sourced threats. This provides security professionals with an entry-level of protection at the perimeters of their infrastructure. However, in this model, once past gateway-based protection and detection technologies, attackers are easily able to move through the flat pools of infrastructure commonly found in cloud data centers. This is primarily due to a lack of internal visibility and control tools that are not present in a Defensive Security model. This is an assumed minimum level of security present in most enterprise security models today.




Key Capabilities

- Perimeter-focused prevention and detection (NGFW, sandboxing, NGIPS, DLP, content security).
- Continuous monitoring of perimeter defenses (e.g. SIEM, sampled network traffic).

Distributed Security

Security professionals who recognize the fundamental imbalance presented by a Defensive Security approach (i.e. “the attacker only has to be right once”) will benefit from adopting a Distributed Security model. A Distributed Security model focuses on controlling risk by gaining visibility and control internally at an intra-workload level by moving protection next to the asset being protected. Once security controls are closest to each workload, architects can then begin to apply these technologies to threat intelligence and the enrichment of log data. When preparing for future levels of multi-cloud security, care should be taken at the Distributed level to select controls that are programmable, scalable, and independent of the workload being protected.



Key Capabilities

- Broad data center visibility (intra-workload).
- Broad controls, regardless of workload placement (intra-workload).
- Programmable architecture (e.g. robust APIs for visibility and control).
- High-efficacy threat intelligence integrated into analytics platforms for the detection of low-sophistication attacks.
- Enrichment of log data with relevant security information.
- Discovery of mode-1 applications destined for cloud-migration.
- Security analytics to consume the data from broad visibility.

Efficient Security

Efficient Security is about making the security system more productive through integration of context and automation to improve the day-to-day security operations of clouds. At this level, security operators can reduce the tedium of security tasks and create a more responsive system in the event of an attack. Security templates become integrated with service catalogs to reduce friction of delivering security controls in a secure, dynamic and intent-driven manner. At this level, the variety of processing actions available expands, including the ability to execute deeper analytics as risk levels increase.




Key Capabilities

- Workload metadata is put into template images for embedding security policy creation during IT-controlled workload lifecycle events from a central service catalog.
- Automated alert triage for high-volume security events.

- Creation of security policy for legacy applications that have been transitioned into the cloud.
- Advanced security analytics combined with deep visibility (e.g. intra-workload packet capture) to provide the visibility required to detect advanced adversaries.

Agile Security

Agile Security increases service creation efficiency by allowing a business application owner (i.e. application developer) to request security policies in a common language of risk. This means that a cloud tenant can request services based upon the needs of an application’s security profile. This federation of policy responsibility requires that a governance model is introduced in tandem so that business application owners are responsible for their policy requests.



Key Capabilities

- Simplified tenant-managed security policy model.
- Cloud governance to provide guard rails for tenant-managed policy.
- Advanced detection and response capabilities (E-W and N-S) in the form of attacker deception techniques and deep packet processing capabilities.
- Conditional policy structure allowing for insertion of appropriate security measures based upon the risk and value of the asset being attacked.
- Increased control depth beyond application-awareness.
- The perimeter security controls have been largely re-architected as services across multi-cloud.

Autonomic Security

Autonomic Security minimizes interaction required for the creation of policies and response to threats – evolving policy definition from declarative creation of relationships by application owners, to the automatic permission of relationships based upon organizational and technical policies (approved application, protocol, and organizational ‘patterns’ and toxic ‘anti-patterns or ‘guide rails’). Machine learning techniques are used to converge policy definitions around observed models and to recognize significant deviations or violations to those models over time, such as those present in a compromised system. At an Autonomic Security level, the risk of deploying workloads and applications to various venues within the multi-cloud is automatically managed by risk brokering system, which evaluates prevailing threat conditions associated with each location, application risk context, and makes deployment decisions accordingly.



Key Capabilities

- Deploy machine learning to further reduce the threat attack surface by adapting to previous attacks and ‘in scope’ changes to application behavior (policy autonomies).

- Predictive analytics of policy definitions based on prior knowledge of service lifecycles and boundaries defined according to corporate ‘acceptable behavior’ models.
- Risk brokerage model informing workload placement decisions in multi-clouds.
- The perimeter has been largely dissolved apart from coarse-grained controls oriented towards the Internet. Security is baked in as part of cloud solution delivery.

Multi-cloud Security

Security for multi-cloud architecture is, by definition, more complex than security for a single cloud. And complexity varies by industry, too – organizations need to comply with regulations and meet requirements such as HIPAA and PCI, which will need to be done differently in multi-cloud environments.



Challenges of Multi-cloud security

- **Data security:** There’s the data itself, potentially spread across multiple providers.
- **Access security:** Then there are the people accessing the data. This is where many companies fail when it comes to security. Employees who shouldn’t have permission to access certain data still do or user accounts aren’t robustly protected, leaving them vulnerable to hackers.
- **Consistent security when scaling:** As you utilize more resources and scale up components, you need to ensure your security is sufficient.
- **Security plan as cloud footprint evolves:** Beyond scaling the resources you’re already using, you’ll likely bring on additional services from new vendors. You need a plan in place to ensure security is fully addressed as each new service is deployed.
- **Synchronizing security policies across vendors:** With each vendor having its own set of controls, it’s challenging to sync decisions across different platforms to ensure consistent policies.
- **Visibility:** Gaining visibility into different platforms, each with its own security features and granularity, is particularly complex in multi-cloud environments.
- **Monitoring:** Each provider offers different monitoring options, but your monitoring must account for the full scope of your entire deployment. Leaving anything out increases the security risk.
- **Increased attack surface:** Multiple providers means a greater number of services and a larger attack surface, giving attackers more ways to infiltrate.



Best Practices for Multi-Cloud Security

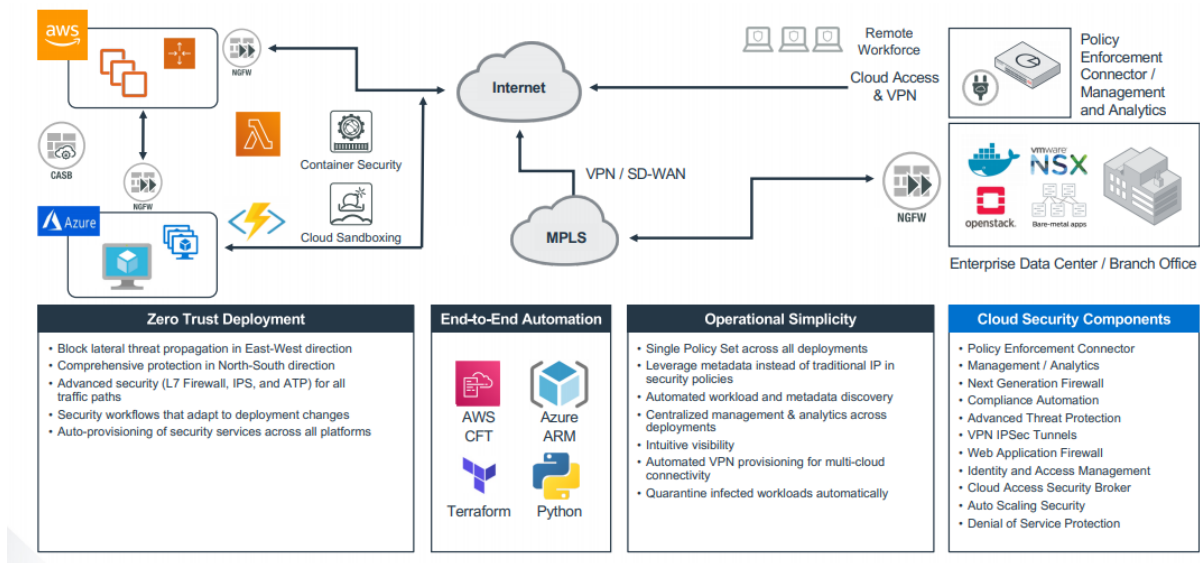


Figure 11 : Security implementation for multi-cloud

Despite these pitfalls, organizations are increasingly choosing the multi-cloud route. With the many advantages multi-cloud offers and the rapid pace of development in the field, it's not hard to see why. By implementing the following best practices, organizations can significantly improve the security of their multi-cloud deployments.

Synchronize policies & settings

If you're using multi-cloud for availability, with identical operations on two clouds, the same security settings should be maintained across both. This can be achieved by synchronizing policies and settings across providers.

Use different security policies for different services

If your organization is using different workloads/applications, individual security policies should be created for each service. For example, if you're planning on setting up a new BI service, the advantages of building it on each platform should be considered first. The security policies should then be based on the chosen platform.

Automate, automate, automate

Using a system that automates various tasks reduces the human risk factor and allows you to stay agile. But be sure to address automation from not only a DevOps perspective but a DevSecOps perspective, to ensure that security is a core consideration and driver throughout the entire process.



Choose the right tools

Find tools and products that allow you to synchronize your security policies across different providers. Your security policies should be written in general terms, with the tools interpreting them based on how your various providers work.

Monitoring

Establish a security monitoring strategy that consolidates logs, alerts and events from different platforms into one location. Tools that automatically remediate issues or provides guidance on remediation strategies are even better.

Compliance

Find tools to help you maintain compliance in a consistent and efficient way across different platforms.



Single point of control

Simplify your sprawl by using a “single-pane-of-glass” tool that gives admins a single point of control to manage all applications and data security across all their cloud deployments.

Minimize “point security solutions”

Minimize the number of “point security solutions,” which don’t integrate well together. Each additional point solution requires expert staff as well as new integrations and deployment. This adds to complexity and increases the likelihood of error.



Similarly, cloud vendors all provide security services. While these may be beneficial within the vendor’s single cloud deployment, they are insufficient when it comes to securing a multi-cloud deployment. You cannot rely on each cloud provider to only protect its own service (for example, AWS to protect your AWS services, Azure to protect Azure, and so on) and assume you’re getting holistic security coverage. You need a single tool that’s capable of providing unified and consistent coverage across all of your deployments.

Multi-cloud security use case “MUSA”

A common reason for not adopting a cloud-based strategy is the perceived lack of security with the claim “cloud is not safe”. This is not true. Security challenges also exist in private cloud environments, or even non cloud-related vulnerabilities such as “weak authentication”.

One of the leading use cases for proposing a Multi-Cloud security framework is the MUSA project. The next section provides an overview of the project’s main sections.

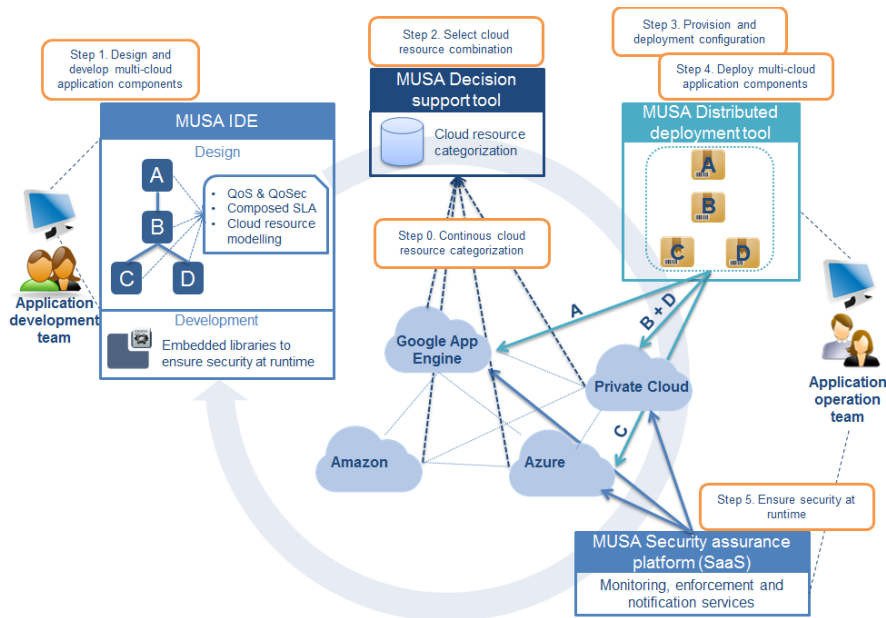


Figure 12: MUSA framework

Main framework components

- A) A preventive security approach, promoting **Security by Design** practices in the development and embedding of security mechanisms in the application.
- B) A reactive security approach, monitoring application runtime to mitigate security incidents, so multi-cloud application providers can be informed and react to them without losing end-user trust in the multi-cloud application.

How MUSA tackles security issues

Figure 15 shows results after the MUSA framework has been applied on a cloud-ready application and the necessary security assessments have been performed.

Security Requirement		DataBase	TSM engine	Identity/Access Manager
	Monitoring	High	High	High
Access Control Requirements	Management Access Control	High	High	High
	User Access Control	High	High	High
Security Requirements Procedure	Countermeasures	High	High	High
	Testing	High	Low	High
	Detection	High	Low	High
	Notification	High	Low	High
	Recovery	High	Low	High
	Key Management	High	High	High

Figure 13: Security assessment results

Design phase

If the security requirement cannot be fulfilled by the cloud provider offering, the missing security capabilities will be addressed by selecting at design phase.

Deployment phase

A set of CSPs that fulfill the best security requirements of the multi-cloud application are recommended by the MUSA Decision Support Tool (DST). Each application component is thus deployed in a CSP that answers the security needs of the application.

Operation phase

The runtime verification enables detection of potential deviations from SLAs and may trigger countermeasures to ensure an optimal security. As an example of a surveyed cloud ready application, a set of security metrics, related to the security controls defined during the risk analysis phase, are captured – as in Figure 15 – using dedicated monitoring agents deployed in the same virtual machine as the application components and set to a centralized MUSA Security Assurance platform.

Defining security SLA

This is done by extracting the Security SLA of the database component of the surveyed application.

```
- <wsag:AgreementOffer xsi:schemaLocation="http://schemas.ggf.org/graap/2007/03/ws-agreement wsag.xsd http://www.specs-project.eu/resources/schemas/xml/control_frameworks/nist nist.xsd">
  <wsag:Name>MUSA_SLA_TEMPLATE</wsag:Name>
  - <wsag:Context>
    <wsag:AgreementInitiator>$SPECS-CUSTOMER</wsag:AgreementInitiator>
    <wsag:AgreementResponder>$SPECS-APPLICATION</wsag:AgreementResponder>
    <wsag:ServiceProvider>AgreementResponder</wsag:ServiceProvider>
    <wsag:ExpirationTime>2014-02-02T06:00:00</wsag:ExpirationTime>
    <wsag:TemplateName>SPECS_TEMPLATE_v1</wsag:TemplateName>
  </wsag:Context>
  - <wsag:Terms>
    - <wsag:All>
      - <wsag:ServiceDescriptionTerm wsag:Name="Database" wsag:ServiceName="storage as service">
        - <specs:serviceDescription>
          - <MUSA:Components>
            - <MUSA:Component name="Database" type="storage as service">
              <MUSA:ComponentProperty name="Database"/>
              <MUSA:Description>Storage of User accounts and User mobility profile</MUSA:Description>
            - <MUSA:Threats>
              <MUSA:Threat name="Missing Function Level Access Control" source="OWASP TOP 10 2013"/>
              <MUSA:Threat name="Unauthorized access to admin interface" source="CSA Survey Cloud Computing Top Threats 2015"/>
              <MUSA:Threat name="Man in the Middle attack" source="CSA Survey Cloud Computing Top Threats 2015"/>
              <MUSA:Threat name="Over-privileged application and accounts" source="CSA Survey Cloud Computing Top Threats 2015"/>
              <MUSA:Threat name="Injection" source="OWASP TOP 10 2013"/>
              <MUSA:Threat name="Sensitive Data Exposure" source="OWASP TOP 10 2013"/>
              <MUSA:Threat name="Data Breaches" source="CSA Survey Cloud Computing Top Threats 2015"/>
              <MUSA:Threat name="Sniffing Storage Traffic" source="Toward a Threat Model for Storage Systems"/>
              <MUSA:Threat name="Exhausting Log, Data and Metadata Space" source="Toward a Threat Model for Storage Systems"/>
              <MUSA:Threat name="Deletion of Data" source="Toward a Threat Model for Storage Systems"/>
              <MUSA:Threat name="Cross-Site Request Forgery (CSRF)" source="OWASP TOP 10 2013"/>
              <MUSA:Threat name="Denial of Service" source="CSA Survey Cloud Computing Top Threats 2015"/>
              <MUSA:Threat name="Modifying Metadata" source="Toward a Threat Model for Storage Systems"/>
            </MUSA:Threats>
          </MUSA:Component>
        </MUSA:Components>
      </MUSA:ServiceDescriptionTerm>
    </wsag:All>
  </wsag:Terms>
</wsag:AgreementOffer>
```

This section gives an overview of the framework and what are the main sections into it.

Security fragmentation

Many companies still have a fractured approach to security. Security technology companies haven't helped this fracturing as most of the approaches out there are to bolt on security features after the fact or even release another product to cover a different aspect of security. Let's be clear, security by its very nature is very hard. There are multiple attack vectors that can lead to almost an infinite amount of different attacks that can come from anywhere. The concept of IoT and Edge computing will fragment and expose even more vectors to attack. Expansion of modern datacenter topology to include multi-cloud, co-locations and edge has simply distributed this fragmentation further out into this new ecosystem.

Security needs more of a systematic approach. The entirety of a system needs to be looked at and the best security practices must be placed at the correct level within the system. While all existing security areas still need to be addressed, in regard to IoT and Edge we have to potentially consider more than just traditional datacenter issues. With the enterprise data center topology becoming more distributed we have all sorts of new issues to deal with. We used to be able to have some physical security constraints, but these are almost all gone when the expansion to the edge happens. Even environmental attributes potentially become new attack vectors.

If we start looking at security in this systematic approach, we can place the best constraints at the best possible location within the system. That is if we don't try to bolt on the security after the fact. We start at the hardware of the system. First, we review the supply chain from which our components come. Trust must be built off this initial building and acquisition. Utilizing secure enclaves at the processor level allows us to start building trust within the system. This trust can also be started by utilizing other trusted hardware with the system.

Once we have a basis of trust the system can start to evolve from a great starting point. Utilizing intelligent connectivity (switches, routers, firewalls, SmartNic, etc.) we can start building out a trusted environment. This includes trusted compute, network, and storage environments. These can't be seen as separate systems. They have to be viewed as a single system. This means both data and control paths have to be simplified and have trust created or transferred throughout the subsystems. This can be accomplished in many ways such as certificates, secrets, etc.

That is how we would currently do this in an enterprise data center. How would this be accomplished when the devices have lower powered processors, installed and initiated by anyone (varying skill set), unreliable, unsecure, or even compromised causes issues in creating that initial trust? There are many other variables to be taken into account; are devices being shared, device never initialized then comes alive after multiple years and so on. The behavior of devices can be altered by environmental and deployment ecosystems. If we start to expand the 'system' view to include these highly distributed edge nodes and devices. This requires these security attributes to be distributed throughout the system again in the best possible location. This might include a local router or an edge gateway.

The emergence of strong privacy protocols, like TLS 1.3, makes traditional monitoring and network snooping almost impossible. The traditional 'man in the middle' monitoring technique will no longer be supported. The only data that will be available will be in the header. Fields such as source, destination,



etc. will be the only information available not encrypted. Inspection of the payload, even with a valid certificate, will no longer be able to decode the message while in transit.

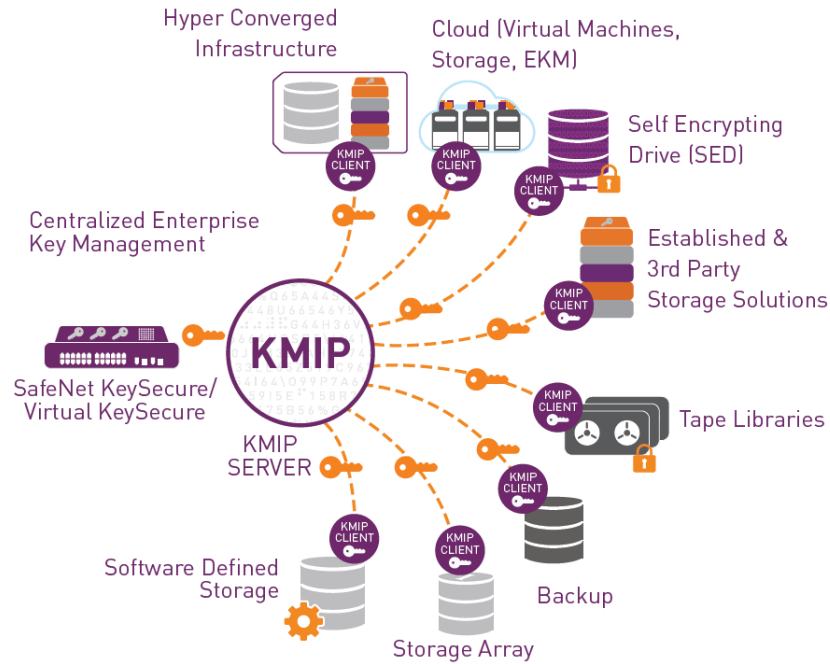


Figure 14:Key management interoperability protocol

Once again, we run into another obstacle trying to distribute secure with these new standards. Certificate management and distribution become an even larger problem in IoT and edge deployments. KMIP and other certificate protocols are adapting to support these new deployment models but there is still a problem with the initial distribution or certificate signing requests (CSRs). Most of these problems occur on devices with limited resources and volatile connections. The amount of data that can be persisted within these devices also cause problems of certificates and validation of certificates include revocation.

IoT and Edge devices also have issues with being properly updated and managed. In some cases, it might not even make sense to enable these management scenarios because of cost or size of the device. Consider the ability to flash the firmware on a sensor that costs \$0.05 USD. Not only are they inexpensive there could be 100s, 1000s, 10000s or even 100000s of these devices deployed in a small area or facilities. Spending time and operational cycles on the devices would drive up the costs so far that it doesn't allow for the usage or deployment of them.

This leads to leaving un-patched and extremely vulnerable code out in the wild with no ability to update them. Even in enterprise data centers today, there is a significant amount of legacy code being executed. Sub-systems or even systems have code that might have been written a decade or more ago for some obscure feature that only one or two customers requested. Leaving this code in updated code basis causes a huge security risk. If this code is now distributed out into this extremely distributed system or ecosystem the risk grows at a rate greater than the number distributed.

Traditionally, to de-risk an enterprise would be to create a 'golden' backup and implement multiple HA and DR strategies. These patterns are hard to achieve in such an environment. The distributed nature of the system and the data created and processed within makes it difficult if not impossible to get a consistent view of the system or data. Additionally, unreliable connectivity of the devices and the inability to rely on them to provide advanced capabilities, i.e. freeze the I/O and then checkpoint it and finally thaw the I/O, is almost impossible. Another consideration is the desire to protect transient data whose value is hard to identify. Missing 10 or 100 or even 1000 data points out of the millions that are generated each day might not make a statistical difference.



In summary creating an extremely distributed system is difficult and completely changes the existing topology of the modern data center, including cloud. Even modern software and hardware development best practices might not fully support distributing infrastructure and data to these extremes. New security standards, protocols, and processes will need to be created in order to help de-risk enterprises looking to expand into these areas. Traditional monitoring and Standard Operating Procedures will need to be revamped to take advantage of the new opportunities that exist out in the edge of computation.

Conclusion

This Knowledge Sharing article illustrates that as more and more companies journey through the digital transformation, they have to take a very systematic approach to security. As the modern data center's topology distributes outside of the traditional boundary to include public clouds, co-location data centers, telco networks and finally some interaction with edge, traditional security patterns and tools will not be adequate. The concept of multi-cloud and democratized access to infrastructure and data will require organizations to upskill their employees with a whole new set of skills.

AI/ML, automation, and policies all play a critical role in the future of highly distributed compute. Creating a secure and trusted distributed compute ecosystem requires more insights and telemetry into somewhat opaque environments today. This systematic approach to security starts at the lowest and farthest-reaching pieces of the system. Enabling trust by knowing how a sub-system boots and its behaviors creates a solid foundation on which to build. Being able to place and automate the correct level of security within the rest of the system is paramount. As we have discussed throughout, the dynamic nature of this new ecosystem requires the security to be just as dynamic and innovative.

Controlling this new ecosystem with the mass number of devices is going to be a difficult task. Though we need to rely on innovative ideas like employing the aka crowd sourcing of the security responsibility to the devices and humans.

Hopefully this article helped prepared you to think differently while driving your organization through its digital transformation.

References

- 1- <http://bit.ly/1K9InEz>
- 2- https://www.bbc.co.uk/news/business-51152151?intlink_from_url=https://www.bbc.co.uk/news/topics/cp3mvpdp1r2t/cyber-attacks&link_location=live-reporting-story
- 3- Figure 1 which presented by FortiGate company as a solution architecture for multi-cloud security.
- 3- <https://www.musa-project.eu/>
- 4- <https://www.continuitycentral.com/index.php/news/technology/4240-traditional-disaster-recovery-vs-draas-six-questions-to-ask>
- 5- https://www.pubnub.com/static/papers/IoT_Security_Whitepaper_Final.pdf
- 6- Wireless sensor network for aircraft health monitoring IEEE paper
- 7- https://www.capgemini.com/resource-file-access/resource/pdf/securing_the_internet_of_things_opportunity_putting_cyber_security_at_the_heart_of_the_iiot.pdf
- 8- <https://enterpriseproject.com/article/2016/2/internet-hackable-things-why-iiot-devices-need-better-security>
- 9- Performance Improvement and Power Consumption Reduction of an Embedded RISC Core, Author: Jung Hongkyun, Jin Xianzhe, Ryoo Kwangki
- 10- Wireless sensor network for aircraft health monitoring, Haowei Bai; M. Atiquzzaman; D. Lilja
- 11- A Trust Evaluation Model for Online Social , IEEE paper
- 12- Networks <https://www.helpnetsecurity.com/2015/04/16/internet-of-everything-attack-surface-grows/>
- 13- <http://www.acpi.info/>
- 14- <https://krebsonsecurity.com/2019/08/what-we-can-learn-from-the-capital-one-hack/>
- 15- https://rhinosecuritylabs.com/aws/capital-one-cloud_breach_s3-cloudgoat/
- 16- <https://towardsdatascience.com/the-real-meaning-and-process-of-data-democratization-abd7d7608a1>
- 17- <https://www.secodis.com/threat-risk-assessments/?lang=en>

Figure 1: Digital transformation	6
Figure 2: Data access Democratization	7
Figure 3: Example of Threat modeling	8
Figure 4: OODA Loop to Accelerate Your Decision Making.....	10
Figure 5: Modern DC Potential attack areas	11
Figure 6: Attack targets	12
Figure 7 : IoT Application Domains & Stakeholders	12

Figure 8: Sensors' behavior algorithm.....	16
Figure 9: Power consumption trust scoring mechanism.....	17
Figure 10 : Securely portioning IoT devices.....	17
Figure 13 : Security implementation for multi-cloud	23
Figure 14: MUSA framework	25
Figure 15: Security assessment results.....	25
Figure 16:Key management interoperability protocol.....	28

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.