Evolve vulnerability management through Cybersecurity Mesh Architecture (CSMA)

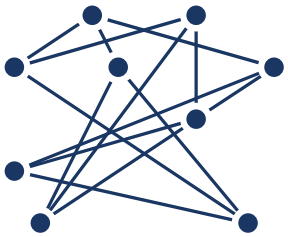

About

Security operations teams must develop an *effective* vulnerability management strategy. A new solution for ensuring effectiveness is through meshing existing data to achieve context. Ensuring interoperability between services provides context to make business-driven security decisions. This emerging concept is often referred to as Cybersecurity Mesh Architecture (CSMA). We outline the strategic goals and benefits of this strategy. Outcomes are effective prioritization for business-specific risk reduction and improved collaboration and communication. For more information on CSMA, please read our whitepaper on the subject.

The founders of appNovi were frustrated by the uncertainty of data they faced as security practitioners in the SOC. Reliance on disparate data sets required tabbing between multiple screens and manual aggregation of data. Convergence and analysis were manual time-consuming processes in Excel. And without complete assurance of the accuracy of data security incident resolution often remained incomplete or indecisive based on ambiguity. This educational guide is to help other practitioners address and eliminate IT data uncertainty using existing network and security tools.

“By 2024, organizations adopting a cybersecurity mesh architecture to integrate security tools to work as a collaborative ecosystem will reduce the financial impact of individual security incidents by an average of 90%.”

Gartner, Top Strategic Technology Trends for 2023: Cybersecurity Mesh

Vulnerability Management

Vulnerability management is the ongoing process of identification, prioritization, and response to risk before exploitation. It is a foundational need for businesses. Effective vulnerability management reduces the likelihood of an attacker accessing their network and limits the pivoting and lateral movements once exploitation occurs.

A vulnerability is a flaw or weakness in system's security procedures, design, implementation, or internal controls. Vulnerabilities require remediation to eliminate risk or mitigation through compensating controls to reduce the probability of exploitation.

Vulnerability management is predicated on the ability to identify vulnerabilities in software and hardware. Software tools such as vulnerability scanners identify vulnerabilities associated with assets and software. Security teams focus on their domain of expertise using different tools, yet specialization results in siloed teams and disparate data sets despite a unified goal.

Vulnerability management is critical for organizations to reduce costs. Patching to prevent exploitation is less costly than responding to a breach based on man hours, recovery, and reputation.

Common challenges in vulnerability management

Too much data and too many types

Vulnerability management professionals must sift through overwhelming amounts of vulnerability information. Each data source has its own format, metrics, method of acquisition, and set of identifiers. These differences make it difficult to correlate and apply a uniform risk policy. Without normalized data, risk can't be understood against the complete data set.

Aggregated data isn't contextual

Vulnerabilities are discovered through agents, scans, network monitors, or external feeds. This results in large data sets siloed across many tools. Analysts must toggle between multiple screens and logins to achieve a common goal as a result. Organizations use a Security Incident and Event Management (SIEM) to aggregate data in a single place. SIEM solutions are the primary data correlation solutions. SIEMs aggregate logs of different tools to refine the alerts into a manageable subset. Yet teams struggle to leverage their data sets for proactive risk assessments. SIEM configurations correlate event data but have limited capabilities to support metadata and telemetry. It is impossible to gain context from SIEM data without manual analysis of different data types and sources. While SIEMs are highly effective data aggregators, they lack the ability to normalize data in a consistent format. The result is often inadequate data to achieve context. Without confidence in data sets, data isn't trusted to initiate security processes. This results in a data lake accessed only after an incident has already occurred. For SIEMs to be effective risk management solutions they need more accurate alert refinement achieved through contextual details.

Risk isn't assessed against business impact

A vulnerability's priority is often determined through a vendor-provided risk score and severity. This objective assessment of risk ignores the significant aspect of context. One hundred assets with the same CVE are weighted equally in importance if we can't determine which are business priorities or those that have compensating controls. Equal severity is assigned to the eCommerce web server and the wirelessly connected espresso machine sharing the same CVE.

Network context is a key aspect to understand the attack surface. Analysts use traffic or security alerts (e.g., IDS) logs, to understand network exposure of the vulnerability and likelihood of an attack. Unfortunately, time-consuming manual processes do not scale with the number of generated alerts.

Without network and business context, vulnerability management isn't aligned to business priorities.

Complexity is the worst enemy of security, and our systems are getting more complex all the time.

Data and Goliath: The Hidden Battles to Collect Your Data and Control Your World
Bruce Schneier

Uncertainty over changes results in inaction

The most challenging aspect of prioritization is understanding the business impact of response. For example, investigating assets by IP addresses requires swivel chair analysis. Analysts tab between many consoles to understand what the asset is beyond its IP address. For example, is it running Linux or Windows? Is it an application server? Does it have security controls in place? Who owns the device? In the case of virtual environments, there may be a question if the asset still exists or has been replaced. With the increased adoption of cloud and emergence of shadow IT traditional asset identification is a more complex challenge. There may be limited data about an asset, or none at all if an employee has signed up for a new service. The inconclusive analysis fails to align an asset to different business processes and units.

The uncertainty of disrupting business processes by making changes that have uncertain results amounts to inaction or partial resolution.

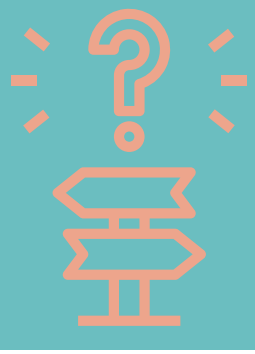

When data for analysis isn't complete, there is uncertainty in the outcomes of response. For example, what applications have a direct dependency on that server? And which applications indirectly depend on that server remaining functional? Understanding a chain of consequences from a single action is necessary to resolve risk. For example, patching a vulnerable database may result in the loss of APIs deprecated by the vendor. If the application server cannot connect to the database server, application functions will fail. And if multiple application servers rely on the database server, they may all be impacted. And we servers that connect to the now broken application servers may not work. One change without understanding context can create a cascade of disruption.

The uncertainty of disrupting business processes by making changes that have uncertain results amounts to inaction or partial resolution.

Why it's time to improve vulnerability management

You have all the data you need, it's just not accessible

All organizations face the same obstacle – data is aggregated but not available in a usable format. When data is not contextually converged analysts must manually retrieve it and converge it. Storing data in a SIEM or data lake creates a repository gated by proprietary query languages, network access, and subject matter expertise. This three-tier gate can prevent access, complicate the methods to acquire data, and result in missed data sources simply because the analyst is unfamiliar with a specific query language. This additional barrier to data requires specialists or costly third-party services. And even with the added layer of enlisted support, subsequent analysis of the returned data requires manual normalization and analysis. Specialists will understand the product and how to query it but may not be familiar with the data and purpose.

Different tools also populate different data consumed by different teams. While application scanners and network scanners are used in parallel, rarely are the two unified into a single data set for analysis. The result is more time spent on data aggregation and analysis.

Risk is always there

There is no future state of a zero-vulnerability network. There will always be more vulnerabilities disclosed while many remain unreviewed or unresolved.

Time spent investigating each alert increases by each data point.
For example, understanding an asset based on IP requires querying several consoles. Understanding network connections requires reviewing NetFlow logs. Understanding the asset's owner is incredibly complex. Analysts query many asset data sources and ask individuals as a result - app owners often change. Escalation results in more time waiting for an answer with the same uncertainty. It also obligates resources that would be more advantageously applied elsewhere.

Vulnerability disclosure is cumulative as thousands of CVEs are disclosed each year. The high frequency of changes to the network creates an increasingly connected and less governed network in parallel. Because there will never be a zero-vulnerability network, prioritization becomes an absolute necessity to manage risk with limited resources. Effective prioritization based on context results in the most risk reduction based on business priority - the preservation of revenue. Yet manual analysis is time consuming to produce meaningful outcomes. And reliance on objective vendor-provided risk scores is reliant on volume and luck.

You cannot analyze all of your existing data without force multiplying the SOC.

Patches don't result in measurable business risk reduction

The measurement of risk is difficult to communicate to executives. Security is generally a technical subject and hard to quantify. As a result, stating the count of remediated and mitigated assets are reported alone does not align to applications dependencies and protected revenue. Translating security decisions to protected revenue is the measurement that executives desire. Tracking remediated assets is not indicative of the business impact of security. Furthermore, it pushes vulnerability teams to pursue quantity over quality. This point is particularly salient when you consider high profile breaches of non-remediated vulnerable web servers with CVEs that were disclosed years prior to exploitation.

Risk isn't accurate when measured in a vacuum

Silos are a sign of strength of a team. Specialists can apply their domain expertise to identify and reduce risk. Specialized tools are employed to implement and maintain security control coverage. The concern is when critical mass of alerts is achieved and there are more tasks than capacity to fulfill. Then risk prioritization becomes more reliant on other factors.

Consider a common example: when every alert is a host with a high severity CVE, all tickets are equal. But not all cyber assets are the same, so this assumption is inherently flawed.

Consider a few answers that require more information than a single tool will provide. Is the asset:

- A workstation or a server?
- Receiving traffic through port and protocol for exploitation?
- Accessible to untrusted networks?
- Subject to anomalous behavior?
- Part of an application if it's a server?
- Indirectly supporting other applications?
- Connected to other assets with similar vulnerabilities?

To answer these questions, we need access to other specialized team's tools and their domain expertise to accurately interpret their data. Gaining access to data is helpful, but understanding it requires specialized knowledge. Engaging other teams is difficult when IT

teams are accustomed to work in silos already with limited resources. Then consider application documentation is an enigma for enterprises at best – applications remain but their developers move on to new projects and application managers change.

Context is the biggest gap to understand and measure risk. This is the significance of the emerging idea of cross domain intelligence (CDI). CDI enables non-specialists to understand specialized subjects through abstraction. Eliminating the reliance on specialists for review makes complex subjects accessible to all.

Security maturity has restarted in cloud-first organizations

Network complexity grows through the introduction and reliance on new technologies. Effective governance is difficult to maintain throughout change. So different teams use different tools to accomplish roughly the same goal. The result is a much more complicated attack surface with disparate data. Meanwhile the distribution of security responsibilities requires re-skilling the newly responsible. But education doesn't make up for experience.

The evidence of this is well supported by the flurry of S3 bucket breaches over the last ten years. Can you recall how difficult it was to get a server to communicate with the internet 15 years ago? It would have been improbable to avoid the security review process. Application owners now bypass many traditional security checks not integrated into the CI/CD pipeline. This isn't a deficiency of developers -- they develop applications and need connectivity. Without appropriate visibility, the dilution of security results in more challenges.

Security must now regain visibility over the enterprise to identify security control gaps and reduce the attack the surface.

Cross domain intelligence enables non-specialists to understand specialized subjects through abstraction. Eliminating the reliance on specialists for review makes complex subjects accessible to all.

Ineffective communication

Security teams already have people, teams, and tools. Reporting is often used by practitioners to report on each team's domain of expertise. Yet technical and granular reporting fails to meet executive requirements for effective consumption. This is especially relevant when considering a security incident.

Motivating individual action often requires escalation into management. Particularly when application owners are assigned that don't have historical context. Security is easily ignored when risk isn't understandable by non-technical stakeholders. Reporting on arbitrary risk scores that reflect technical details and lack business context are reports that security teams will use but others will have difficulty in consuming.

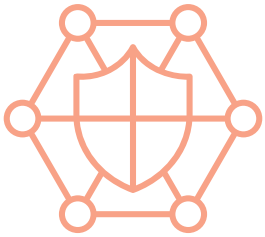

Meshing existing tools and data for effective vulnerability management

Relying on a single data point is an example of a constructed information silo. We have multiple network and security data sets created by tools. The historical challenge has been converging data points into an accessible format. Now integrations and machine learning can eliminate the uncertainty of data and ensure accessibility. Accessibility to security provides:

- Demonstrable and increased ROI from existing tools
- Upleveling of junior analysts to address skill barriers
- More effective communication of security
- Shortened incident response time

Let's consider data that every organization has data for:

1. **Assets** – Stored in your CMDB(s), IaaS solutions, EDR, vulnerability scanners, and inventory tools
2. **Vulnerability** – Often produced by vulnerability scanners, vulnerability correlation tools, or endpoint agents
3. **NetFlow** – Generated by network virtualization technologies. Also available in your L2, L3, and L4 devices or commercial products.
4. **Alerts** - Generated by IDS and IPS sensors or other monitoring solutions aggregated in SIEMs
5. **Identity** – Available through third parties or network virtualization technologies
6. **Application data** – At least the clues to gain this. Understanding application servers and their connections provide the ability to discover applications. There are third party tools for this, but few with successful implementation.

Each data point represents an aspect of risk. Through data convergence, you gain context. The more data points, the better the context. The better the context, the more informed your perspective is. Hence enterprises integrating security tools into a cooperative ecosystem. When composable and scalable, this aligns to Gartner's cybersecurity mesh architecture approach. An ideal framework with the goal of reducing costs of breaches by 90%.

Users, devices, applications and data have left the traditional office and data center. This means that a single network perimeter no longer exists to judge that "inside is good and outside is bad." Identity and context have become the ultimate control plane in a distributed environment that supports assets and access from everywhere.

Top Strategic Technology Trends for 2023: Cybersecurity Mesh
Gartner

Achieving context through cybersecurity meshing of data

You gain a high fidelity understanding of your network and security when data meshes. Consider the following outcomes of combing the above examples of common data sources.

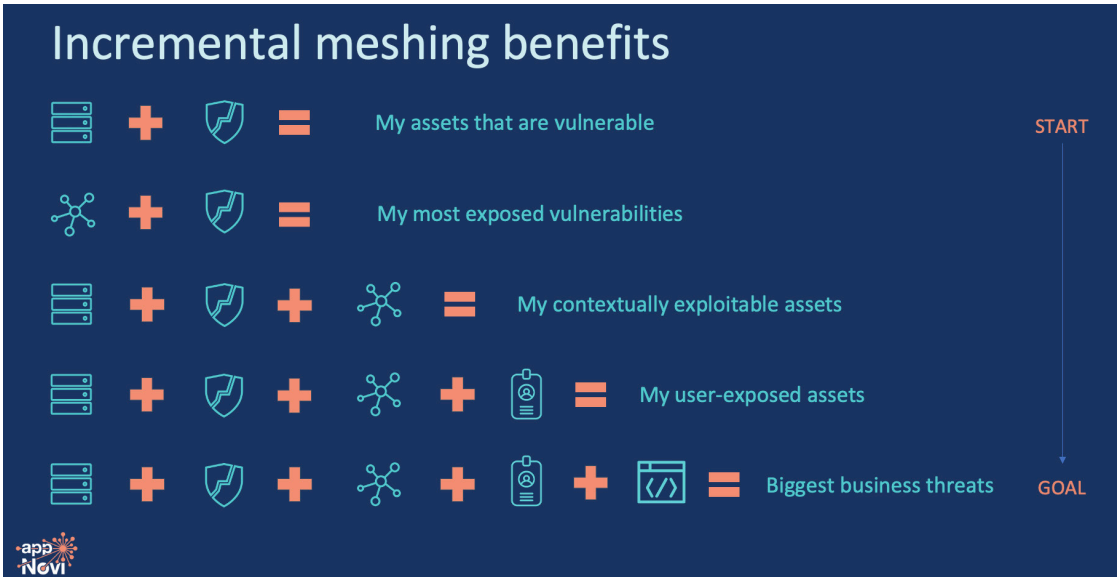

Asset + Vulnerability Data

My assets with vulnerabilities. This helps me determine the vulnerable assets. The problem is that everything is vulnerable.

Vulnerability + NetFlow Data

Vulnerabilities accessible for exploitation through existing access. I know the most relevant patches to my organization based on existing connections. I can choose to patch based on the most risk reduction based on network context.

Asset + Vulnerability + NetFlow Data

My vulnerable assets exposed through network connections for exploitation. I know which assets are contextually exploitable by untrusted networks. I also know indirectly connected assets on my network to understand the blast radius. If I include events, I know which locally exploitable vulnerabilities are most targeted.

Asset + Vulnerability + NetFlow + Identity Data

When I mesh identity events with my cyber assets and connections, I identify insider threats based on behavior. I see multiple failed logins. I see outbound requests to untrusted networks from vulnerable assets. I automatically identify admins with domain rights without MFA enabled. I understand who owns an application server by from admin logins for incident response.

Asset + Vulnerability + NetFlow + Identity Data + Application Data

When I understand which servers are application dependencies, I can track applications. I understand which assets are indirect dependencies of other applications. I understand if making changes during a security response will break application connections. Security now works with assurance that we are protecting revenue.

Understanding the difference between a spark and an inferno

Fighting fires

Fire follows paths of least resistance. The path it follows is based on multiple factors. There are flammable materials. Chimneys inadvertently funnel and amplify heat (i.e., why laundry chutes are illegal in most houses). Rooms without fire mitigation solutions like sprinklers are more likely to burn. Environmental influences influence the reach and intensity of flames across a structure and other houses or may result in natural smoldering. All houses are flammable but not equally.

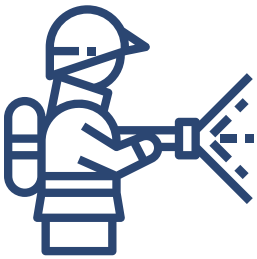

Hackers, like fire, seek the path of least resistance. We have a problem when we don't have context to understand the paths.

Let's consider the ideal scenario when we have context of a house. There's a report of a fire, and everyone we need is there. The firemen get blueprints from the city archive to understand origination points. The builder knows the flammable materials to predict the paths of least resistance. The structural engineer knows what's most critical for preserving to reduce reconstruction costs. The meteorologist tells us the wind strength and direction, so we know what to defend. And of course, the chief coordinates firefighters to extinguish flames. When we have context, we save as much as we can and prevent bad from getting worse.

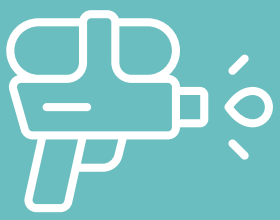

Without context we lack efficiency and assurance of outcomes. Our only solution is to throw more resources at the problem. Hire more firefighters, send more trucks, call in air support to dump water from above... we can only escalate resources based on availability. But we don't have prioritization or efficiency. Without access to data to develop a strategy and an effective response we rely on luck. Without direction, our hoses may be as effective as squirt guns. Every spark, whether a flicker or an inferno, is treated equally based on our risk tolerance. And there's always a new flickering.

Understanding in reality

The sudden disclosure of Log4j sent security teams and software providers scrambling. Many organizations found themselves in the middle of a proverbial wildfire. SOC analysts had longer shifts. Applications teams huddled down. Business leaders coordinated the communication and response across teams and to the public. The disclosure was an unwelcome test of response plans. And the prevalence of the vulnerability alone triggered a massive prioritization effort.

Security teams all faced the same question. “When everything is vulnerable, what are the most important assets to secure?”

Too often we rely on response volume to reduce risk. We gain a mix of wins and losses where the impact of each outcome is dependent on luck. Perhaps most frustrating is the immeasurable business impact.

Successes in firefighting

New vulnerabilities are always found. Known ones are targeted by attackers. It’s imperative to develop effective vulnerability management and response to new disclosures. We know threats spark across the network. Some fizzle out and others have a higher probability to grow. The difference between them is context. We need to understand the business impact of threats. We have our network, security, and application data -- we need to make it accessible to our teams.

Consider the prior example of Log4j. We need to understand which vulnerable assets connect to the internet. We need to know which of those are application servers. What other vulnerable assets connect to the vulnerable and internet exposed assets? Who owns those machines? Can you apply compensating controls without disrupting applications? Can you quarantine an asset without network disruption? We determine priority and available actions when we have this information. We know who we need to work with to remediate the assets. We abstract the issue into a visualized format for technical and non-technical stakeholders.

Too often we rely on response volume to reduce risk. We gain a mix of wins and losses where the impact of each outcome is dependent on luck. Perhaps most frustrating is the immeasurable business impact.

Ransomware attacks are simple. The outcomes are complex. Prevention needs to be simple too.

Reliance on single data points increases complexity. Meshing other data points simplify triaging. Take the data associated with identity for example. Identity is often not a part of vulnerability management. Yet most successful ransomware attacks start with it. Active Directory admin accounts without MFA turned on are ideal targets for credential stuffing. Passwords not rotated through a Privileged Access Management (PAM) system are prime candidates. You already have the data points from Azure (for example). You can alert on these or even automate a check against leaked password repos for a match. You can increase priority for locally exploitable assets in parallel as they are at increased risk.

Maintaining automated analysis across a converged data set provides immediate material benefits. Risk management and incident response teams' benefit, and your cyber insurance underwriter.

Understanding business importance needs to be immediate

CVE scores are helpful for tracking vulnerabilities across cyber assets, but each vulnerability requires a unique set of circumstances for exploitation. Some may require a local action by an asset user, others simply require a port and protocol for exploitation. Alternatively, some require physical access to a device. CVE severity scores are helpful to understand the ease by which an asset could be exploited but cannot be relied on to tell you the exploitability of an asset in the current context. Deciding which patches are selected, or if an asset is part of an application, requires context. For example, consider a prevalent Linux vulnerability is associated with ten cyber assets but none of them are dependencies of an application.

Each vulnerability requires a unique set of circumstances for exploitation... CVE severity scores are helpful to understand the ease by which an asset could be exploited but cannot tell you the exploitability of an asset in the current context.

An equally risky SQL vulnerability is only associated with two cyber assets, but one is an application server located in the same subnet as the other which receives network traffic from contractors. Objectively resolving ten high risk assets sounds more attractive than two but remediating the two assets reduces the most amount of risk to the business. This level of understanding is necessary to make the most risk reduction from the business perspective but supporting this analysis with manual efforts is an impossibility.

Cross Domain Intelligence (CDI) is the transfer of specialized technical knowledge. It enables quick understanding of complex subjects without requiring SME involvement. Security specialists face significant challenges understanding risk and communicating it without business context. When specialized data abstracts into normalized, searchable, and intuitive data, CDI is established. Identifying, prioritizing, and responding to risk is much faster with context. Non-technical or specialized stakeholders can understand complex subjects. The outcome is more effective security communication aligning risk to revenue protection.

Fixing the banality of scan and patch

The reality is no single vendor provides a risk score that is unique to your business

Anton Chuvkin wrote "We scan, we patch, but we don't do vulnerability management" while at Gartner. He noted that vulnerability management teams focused on business risk reduction just patch faster. This method focuses on volume to patch the largest count of vulnerable machines. Or patching the most severe vulnerabilities, or a combination of the two.

The reality is no single vendor provides a risk score that is unique to your business. Vendor-provided scores are unique to the method and type of data they assess. Implementing a risk-based vulnerability management strategy requires context gained from all data sources. Convergence eliminates ambiguity and enables action. Analysts are more efficient and effective when data is at their fingertips. Patching relies on business impact and network context for priority. Outcomes are measurable through revenue protection. Processes are repeatable and analysis is automated. The result is a more efficient SOC and increased availability of specialized resources.

Patching relies on business impact and network context for priority. Outcomes are measurable through revenue protection.

A picture is worth ten thousand rows

Manual aggregation and interrogation of data is time-consuming. It was our biggest pain point working in the SOC. Bouncing between ServiceNow, Tenable, Panorama, writing Splunk queries, contacting server owners... hours would be spent on one alert. And we would be frustrated to refine the data to find the box is actually... the cafeteria menu server.

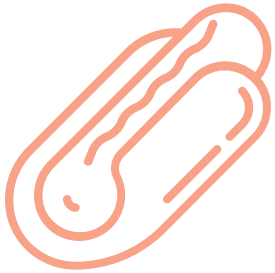

That's why we founded appNovi.

Effective data convergence, machine learning, and modern graph software eliminate manual efforts. Visual presentation of collapsed data sets provides focused investigations. Complex subjects abstract into a simple form and are accessible to all. Network connections and security events are understood. We need to make security more accessible. When we do that, we uplevel junior analysts. We automate security tedium. We empower informed non-disruptive decision-making by eliminating IT uncertainty. We make sure your data is accessible to all stakeholders and not behind a query language gate.

And we make it an experience. Something you can understand, your boss can understand, their boss, and all the way up. We made security accessible so security teams don't waste their time like we did. We were in those seats. We had the same frustration. Now we have a solution.

You'll never have 0 vulnerabilities. You will always have data. Now you can make it manageable and uplevel your SOC.

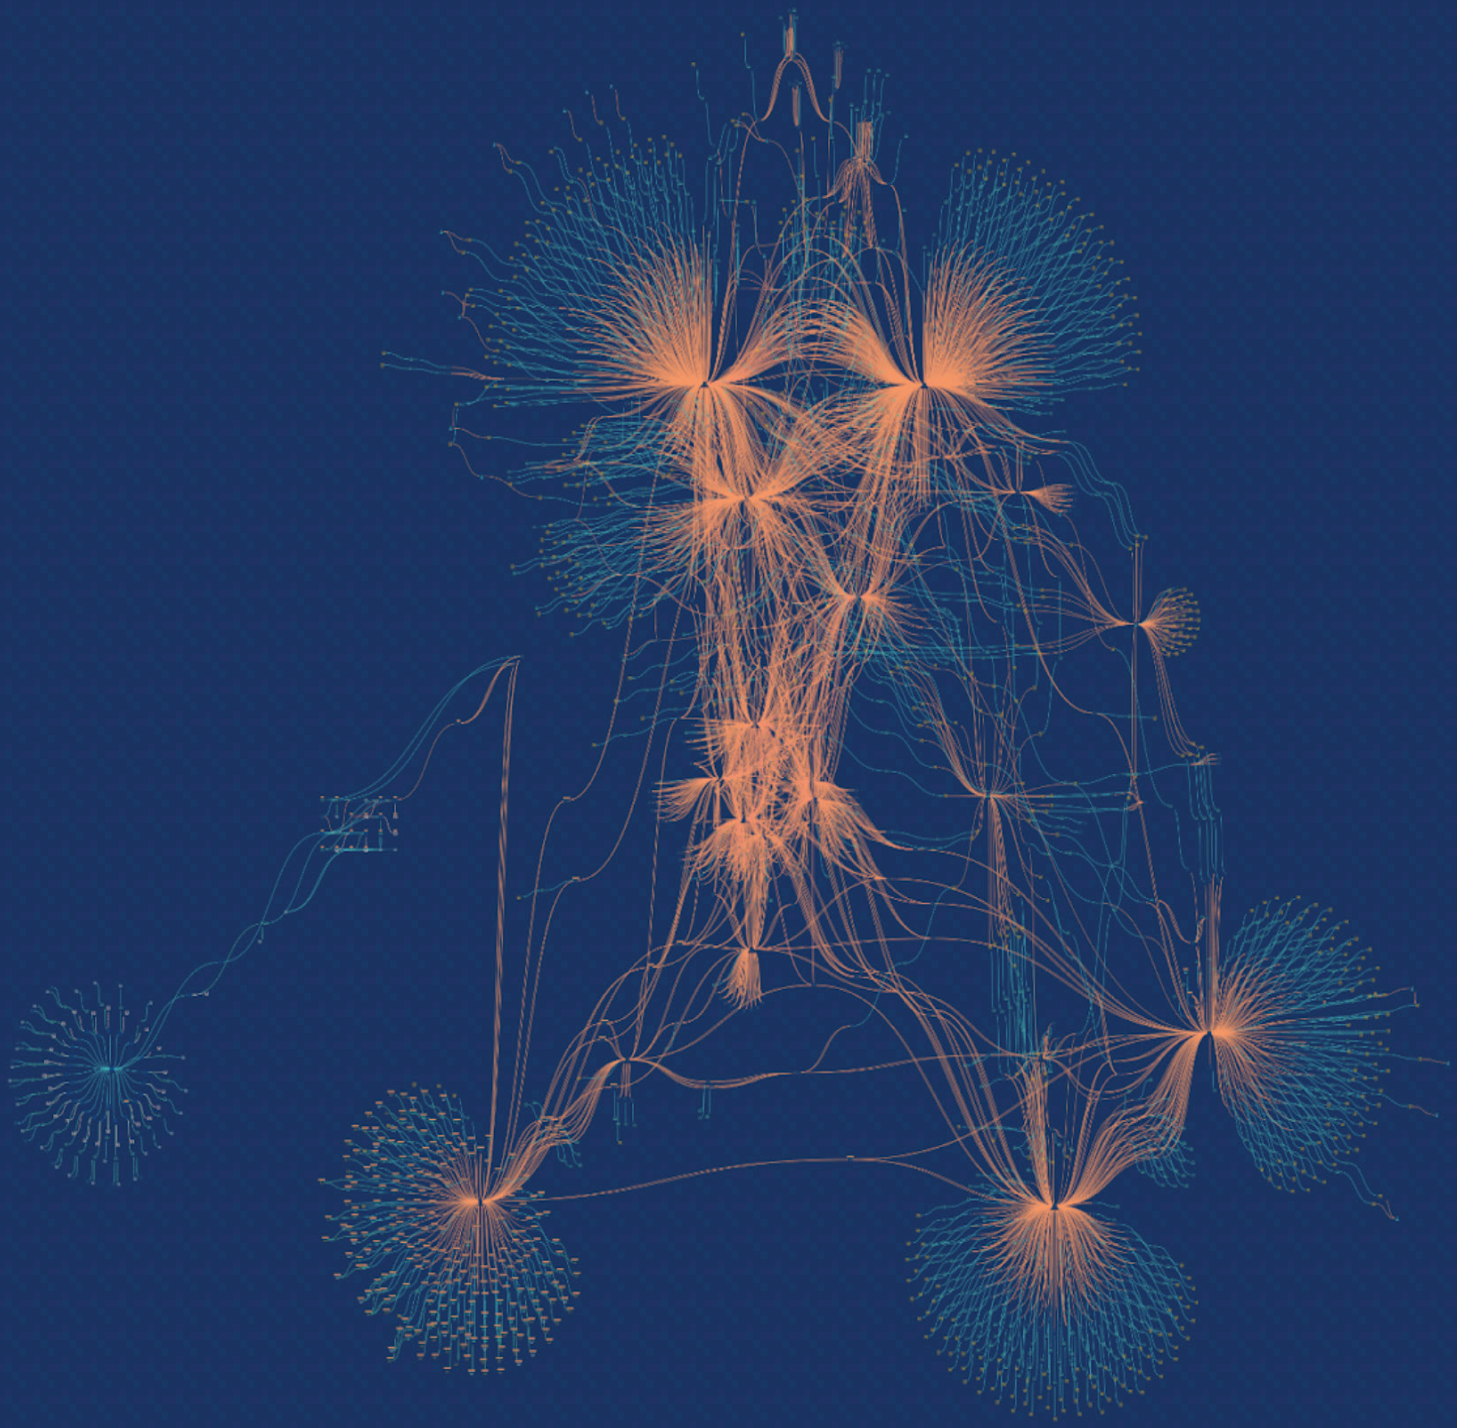

About appNovi

appNovi provides cybersecurity mesh architecture by integrating with your existing network and security services. appNovi customers mesh and visualize their data for network-wide attack surface mapping, vulnerability prioritization based on business risk, and enable efficient non-disruptive incident response.

Gain total asset visibility