EAG | EXECUTIVE ADVISORS GROUP

SD

Reasonable Security Implications of the TRUST Ordinance

-Surveillance Impact Report (§210.0102 (p) (1) – (12))

-Surveillance Use Policy (§210.0102 (q) (1) – (11))

Matt Stamper, PAB Security Professional
CEO | CISO Advisor, Executive Advisors Group, LLC
https://www.linkedin.com/in/stamper/

Reasonable Security
A CISO's Definition

"Reasonable security is that level of security capability that meets the organization's agreed-to risk tolerances while fulfilling regulatory requirements and contractual obligations of the organization."

Reasonable Security
A CISO's Definition

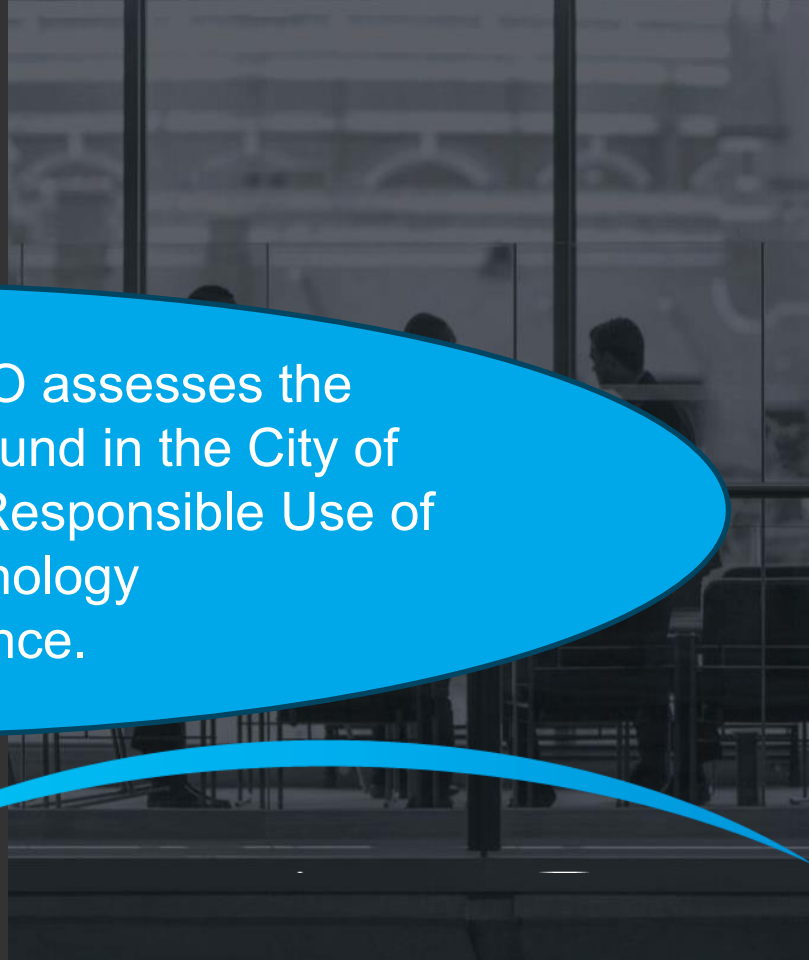

Perspectives on how a CISO assesses the cybersecurity requirements found in the City of San Diego's Transparent and Responsible Use of Surveillance Technology (TRUST) ordinance.

Surveillance Impact Report

Surveillance Impact Report (§210.0102 (p) (1) – (12))

The Surveillance Impact Report requires details on the following:

- Description (p)(1)
- Purpose (p)(2)
- Location (p)(3)
- Impact (p)(4)
- Mitigation (p)(5)
- ***Data Types and Sources (p)(6)***
- ***Data Security (p)(7)***
- Fiscal Costs (p)(8)
- ***Third Party Dependence (p)(9)***
- Alternatives (p)(10)
- Track Record (p)(11)
- Public Engagement and Comments (p)(12)

Surveillance Impact Report (§210.0102 (p) (1) – (12))

Data Types and Sources (p)(6)

Data Types and Sources: A list of all types and sources of data to be collected, analyzed, or processed by the surveillance technology, including scores, reports, the logic or algorithm used, and any additional information derived from the surveillance technology, except that no confidential or sensitive information should be disclosed that would violate any applicable law or would undermine the legitimate security interests of the City.

Surveillance Impact Report (§210.0102 (p) (1) – (12))

Data Types and

Da

s

be d

legitimate s

Reasonable security questions for the use of surveillance technology could include:

- Is there an accurate and complete data inventory?
- Is there an accurate and complete inventory of data sources?
- How are data elements labeled?
- Are data elements co-mingled across other systems?
- How are algorithms assessed for their accuracy?
- Do the algorithms introduce any bias or unintended impacts?

Surveillance Impact Report (§210.0102 (p) (1) – (12))

Data Security (p)(7)

Data Security: Information about the controls that will be designed and implemented to safeguard the data collected or generated by the surveillance technology from unauthorized access or disclosure, except that no confidential or sensitive information should be disclosed that would violate any applicable law or would undermine the legitimate security interests of the City.

Surveillance Impact Report (§210.0102 (p) (1) – (12))

Data Security (

Da ted

in

unde

Reasonable security questions for the use of surveillance technology could include:

- Which security controls have been implemented?
- How are these controls assessed for both design and operating effectiveness?
- How frequently are the controls assessed?
- Are there gaps in required controls, and what is the plan to address these?
- Which control standard or framework is being used?

Surveillance Impact Report (§210.0102 (p) (1) – (12))

Third Party Dependence (p)(9)

Third Party Dependence: Whether use or maintenance of the surveillance technology will require data gathered by the surveillance technology to be handled or stored by a third-party vendor at any time.

Surveillance Impact Report (§210.0102 (p) (1) – (12))

Third Party Dep

Thi

Reasonable security questions for the use of surveillance technology could include:

- Which services are provided by third parties?
- What are the CUECs required of the City of San Diego to use the technology?
- How were the third-party's security, privacy, availability, and governance controls assessed? How frequently are these controls re-assessed?
- Is the third-party required to notify the City of a data breach?

Surveillance Use Policy

Surveillance Use Policy (§210.0102 (q) (1) – (11))

The Surveillance Use Policy requires details on the following:

- Purpose (q)(1)
- Use (q)(2)
- ***Data Collection (q)(3)***
- ***Data Access (q)(4)***
- ***Data Protection (q)(5)***
- ***Data Retention (q)(6)***
- Public Access (q)(7)
- ***Third Party Data Sharing (q)(8)***
- Training (q)(9)
- ***Auditing and Oversight (q)(10)***
- ***Maintenance (q)(11)***

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Data Collection (q)(3)

Data Collection: The information that can be collected, captured, recorded, intercepted, or retained by the surveillance technology, data that may be inadvertently collected during the authorized uses of the surveillance technology and what measures will be taken to minimize and delete the data, and any data sources the surveillance technology will rely upon, as applicable, except that no confidential or sensitive information should be disclosed that would violate any applicable law or would undermine the legitimate security interests of the City.

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Data Collection

Da

a

sour

confidentia

applicable law or wou

ate any

ests of the City.

Reasonable security questions for the use of surveillance technology could include:

- What are the data governance practices associated with the surveillance technology?
- What ensures that only authorized data is collected?
- How is sensitive data protected and governed through its lifecycle?
- How are cryptographic services used with the data?
- How are CRUD actions governed (e.g., access, logging, etc.)?

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Data Access (q)(4)

Data Access: The job classification of individuals who can access or use the collected information, and the rules and processes required prior to access or use of the information, except that no confidential or sensitive information should be disclosed that would violate any applicable law or would undermine the legitimate security interests of the City.

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Data Access (q)

Da

a

secu

Reasonable security questions for the use of surveillance technology could include:

- Who is authorized to access the data?
- How is this access logged and secured?
- How frequently are access and entitlement reviews performed?

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Data Protection (q)(5)

Data Protection: The safeguards that protect information from unauthorized access, including system logging, encryption, and access control mechanisms, except that no confidential or sensitive information should be disclosed that would violate any applicable law or would undermine the legitimate security interests of the City.

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Data Protectio

Da ss,

a

Reasonable security questions for the use of surveillance technology could include:

- What type of authentication is used when accessing the data?
- Which technical, administrative, and physical controls are used?
- How are these controls evaluated for both design and operating effectiveness?
- How frequently are these controls assessed?
- Are there gaps in the controls, and what is the plan to adequately address these gaps?

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Data Retention (q)(6)

Data Retention: The time period, if any, for which information collected by the surveillance technology will be routinely retained, the reason the retention period is appropriate to further the purposes, the process by which the information is regularly deleted after that period lapses, and the specific conditions that must be met to retain information beyond that period.

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Data Retention /

Da

re

met

Reasonable security questions for the use of surveillance technology could include:

- How long is the data retained and what happens at expiry?
- In which jurisdiction or locale is the data retained?
- What type of media is used to back up the data?
- Who has access to the backup data?
- How is this access validated?
- Are there third parties required to support data retention, and have they been assessed and their security controls validated?

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Third Party Data Sharing (q)(8)

Third Party Data Sharing: If and how information obtained from the surveillance technology can be accessed or used, including any required justification or legal standard necessary to do so and any obligations imposed on the recipient of the information.

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Third Party Dat

Thi

in

Reasonable security questions for the use of surveillance technology could include:

- Which third parties have access to the surveillance data?
- Is that access authorized?
- How is the data accessed (e.g., API, direct connection to the database, reports, administrative access to the system, etc.)?
- Which security controls are in place over third-party data sharing?
- Is the data encrypted in transit and at rest with the third party?
- Does the third party share the data with other entities?

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Auditing and Oversight (q)(10)

Auditing and Oversight: The procedures used to ensure that the Surveillance Use Policy is followed, including identification of internal personnel assigned to ensure compliance with the policy, internal recordkeeping of the use of the surveillance technology and access to information collected by the surveillance technology, technical measures to monitor for misuse, identification of any independent person or entity with oversight authority, and the legally enforceable sanctions for violations of the policy.

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Auditing and Ov

Audi

P

te

or en

of the poli

ons

Reasonable security questions for the use of surveillance technology could include:

- What is the nature and extent of access reviews over the surveillance technology and its data?
- How are personnel trained on requisite privacy, security, and governance requirements?
- What type of logs and monitoring are implemented over surveillance technology and access to the technology's data?
- How are third parties reviewed and assessed?
- Which oversight procedures are managed centrally by the City versus the City Department or Agency?

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Maintenance (q)(11)

Maintenance: The procedures used to ensure that the security and integrity of the surveillance technology and collected information will be maintained.

Surveillance Use Policy (§210.0102 (q) (1) – (11))

Maintenance (

M

Reasonable security questions for the use of surveillance technology could include:

- How is the technology maintained?
- What type of access is required to maintain the technology?
- Who is responsible for ensuring that the technology is patched?
- How frequently are patches implemented?
- Have SLAs/SLOs been established, and are they consistent with RPO/RTO requirements for the City of San Diego?

AI is Increasingly Used With Surveillance Technology

The basic due diligence in the preceding slides only touches the surface of how IT and cybersecurity practices are evolving based on the adoption of AI across organizations and how AI is embedded in technology.

Recent news events are a wake-up call for IT and security professionals, both within the City of San Diego and for the organizations supporting City services, including this Privacy Advisory Board. Key examples include:

- JADEPUFFER
- Hugging Face
- Mythos (Anthropic) & Daybreak (OpenAI)


Fascinated by Reasonable Security?

Reasonable Security is Increasingly Mandated

- The California Consumer Privacy Act (CCPA)/California Privacy Rights Act (CPRA) require 'reasonable' security - CIV § 1798.81.5 (see - https://leginfo.legislature.ca.gov/faces/codes_displaySection.xhtml?lawCode=CIV&sectionNum=1798.81.5)
- The Defense Federal Acquisition Regulation Supplement (DFARS) requires adequate security over Controlled, Unclassified Information (CUI) and Federal Contract Information (FCI) (see sections 252.204-7012 & 252.204.7020) and requires covered entities to fulfill the controls found in NIST 800-171 and NIST 800-171a.
- The Securities and Exchange Commission (SEC) requires that public entities disclose material cyber incidents within 4 days of the materiality determination and status regarding the nature and extent of their security programs (see - https://www.sec.gov/files/rules/final/2023/33-11216.pdf).

Reasonable Security
A CISO's Definition

"Reasonable security is that level of security capability that meets the organization's agreed-to risk tolerances while fulfilling regulatory requirements and contractual obligations of the organization."


There’s a Reason Why Security is Tough!

Risk is...

Security is challenging

…

and there is no perfect security.

Good security programs still face material risks.

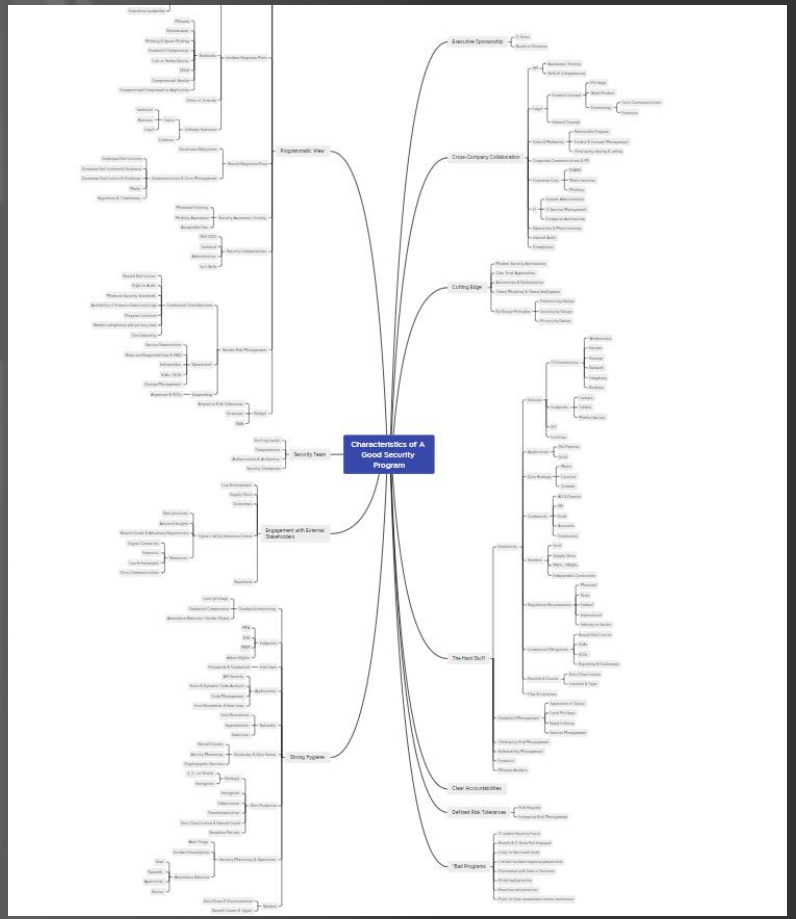

Why security is tough...

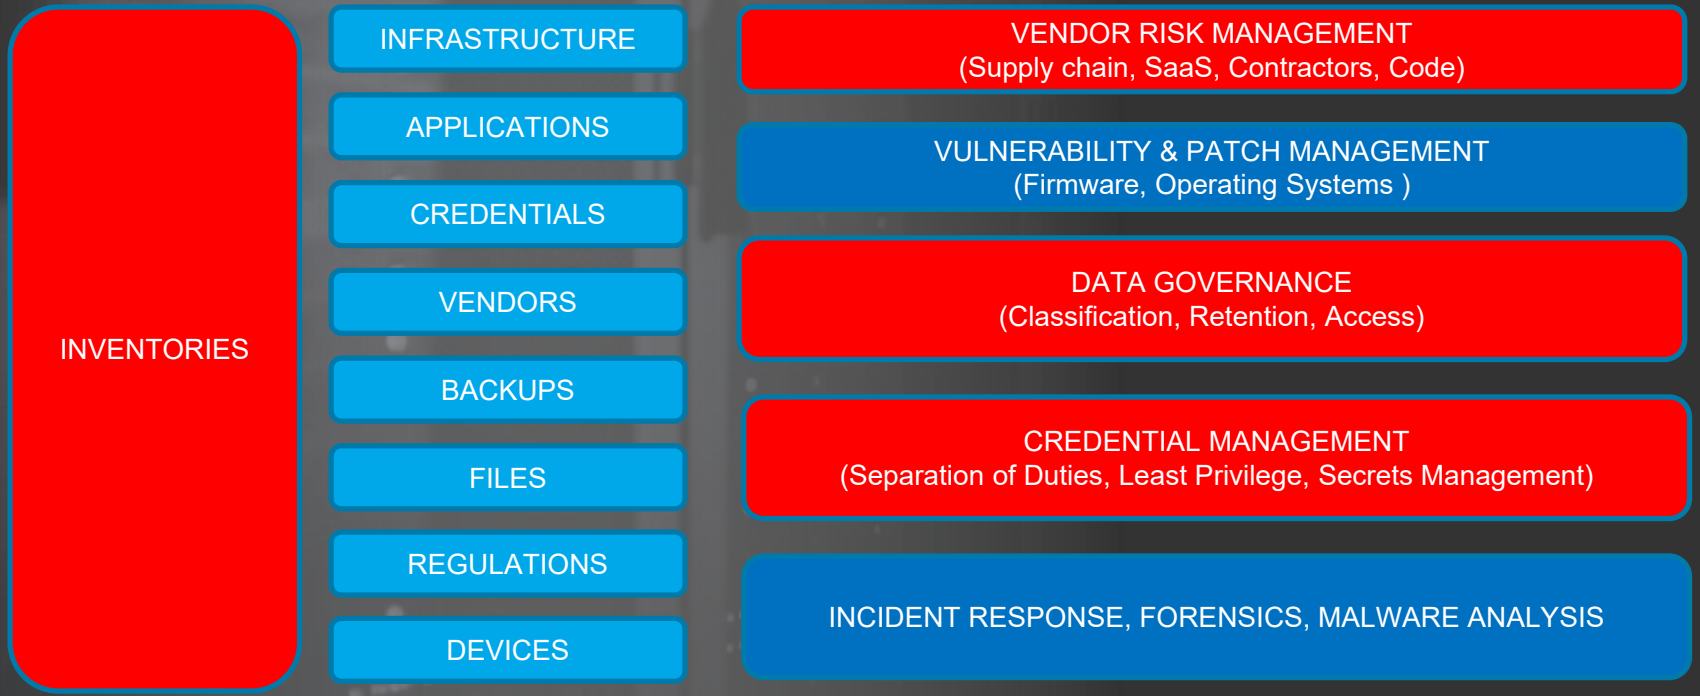

Why security is tough...

How have these practices changed with the adoption of AI tools and applications in our operating environment?

Why security is tough...

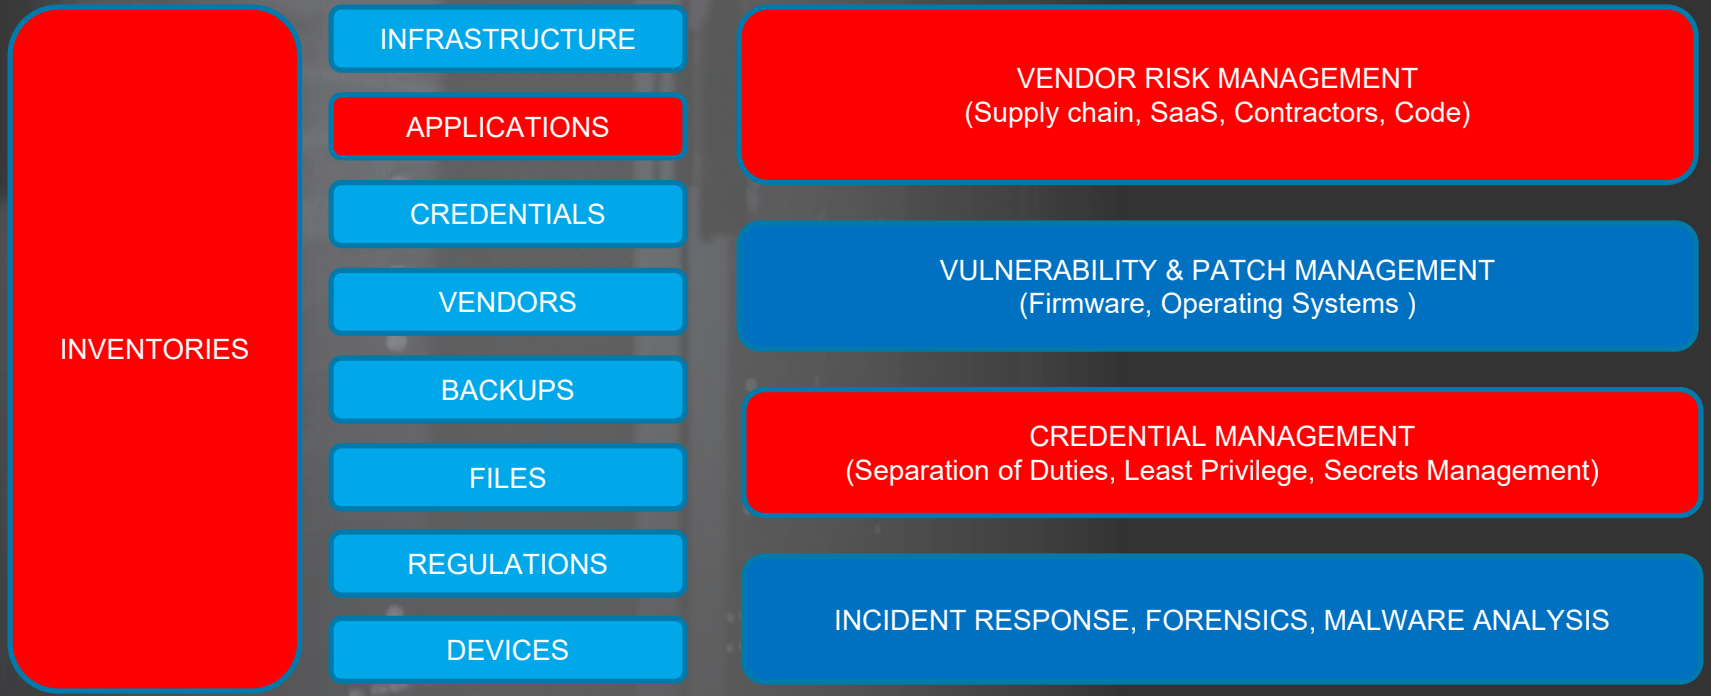

Why security is tough...

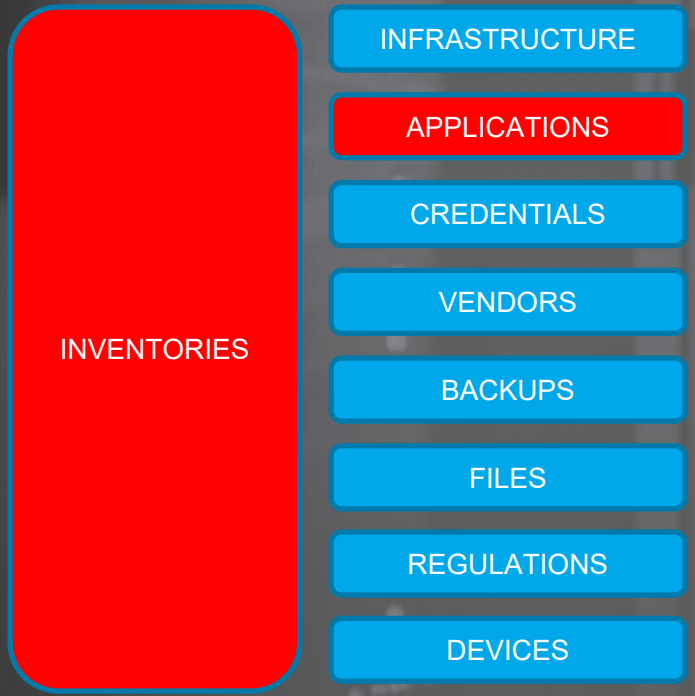

Should we think of AI applications as distinct from other third-party applications?

- User permissions and entitlements?
- Application validation and assessment (think OWASP ASVS)?
- Application code and the provenance of the contributors?
- APIs and integrations between and among applications (think MCP)?
- Data governance?
- Application reviews and assessments?
- Complementary User Entity Controls (CUECs)?

Why security is tough...

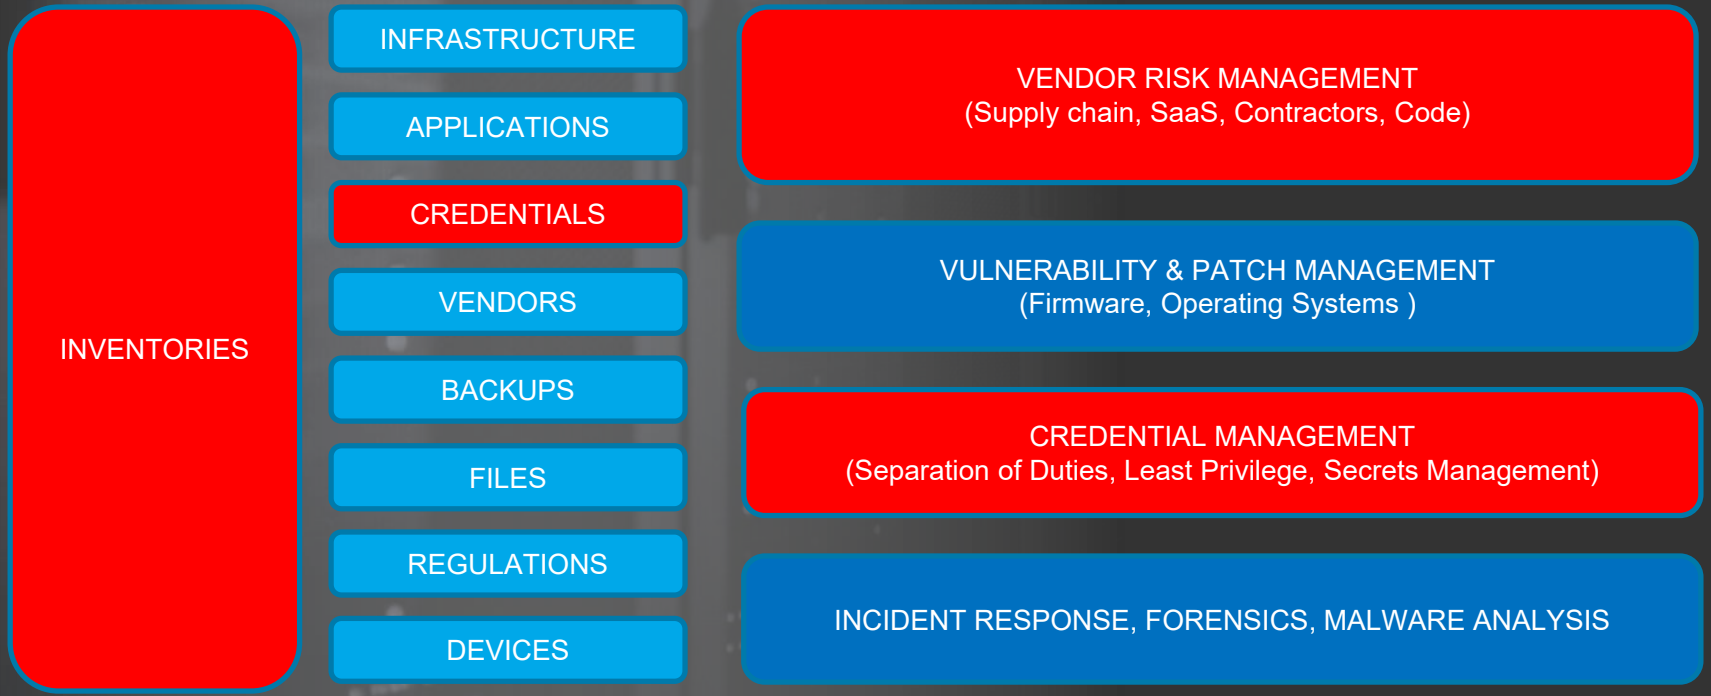

Why security is tough...

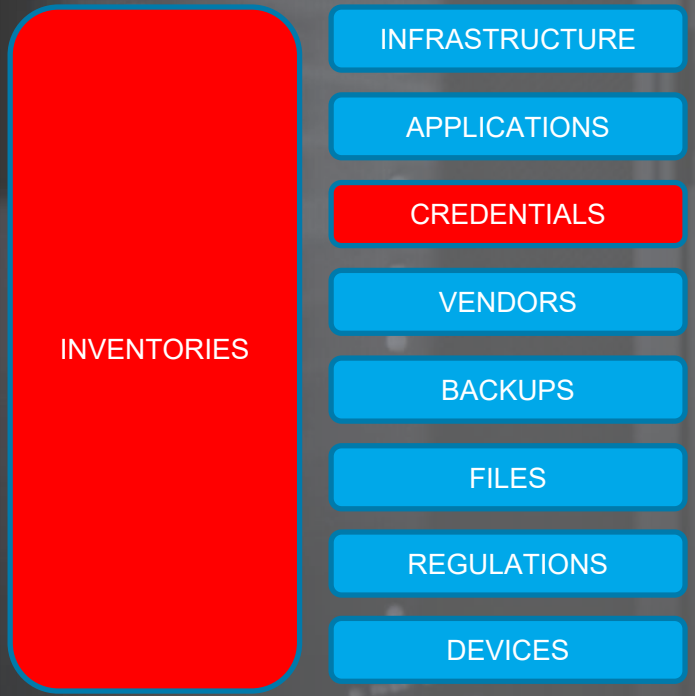

There are critical considerations with respect to identity governance and AI, notably agentic AI, including the following:

- What are the entitlements given to agents?
- Were these entitlements authorized?
- How are the agents' actions and decisions monitored for appropriateness?
- What happens if an agent goes rogue?
- What happens if an 'agent' account is taken over?
- How frequently should entitlement reviews be performed?
- What are the 'on-boarding and off-boarding' procedures for an agent?

But good security is achievable…

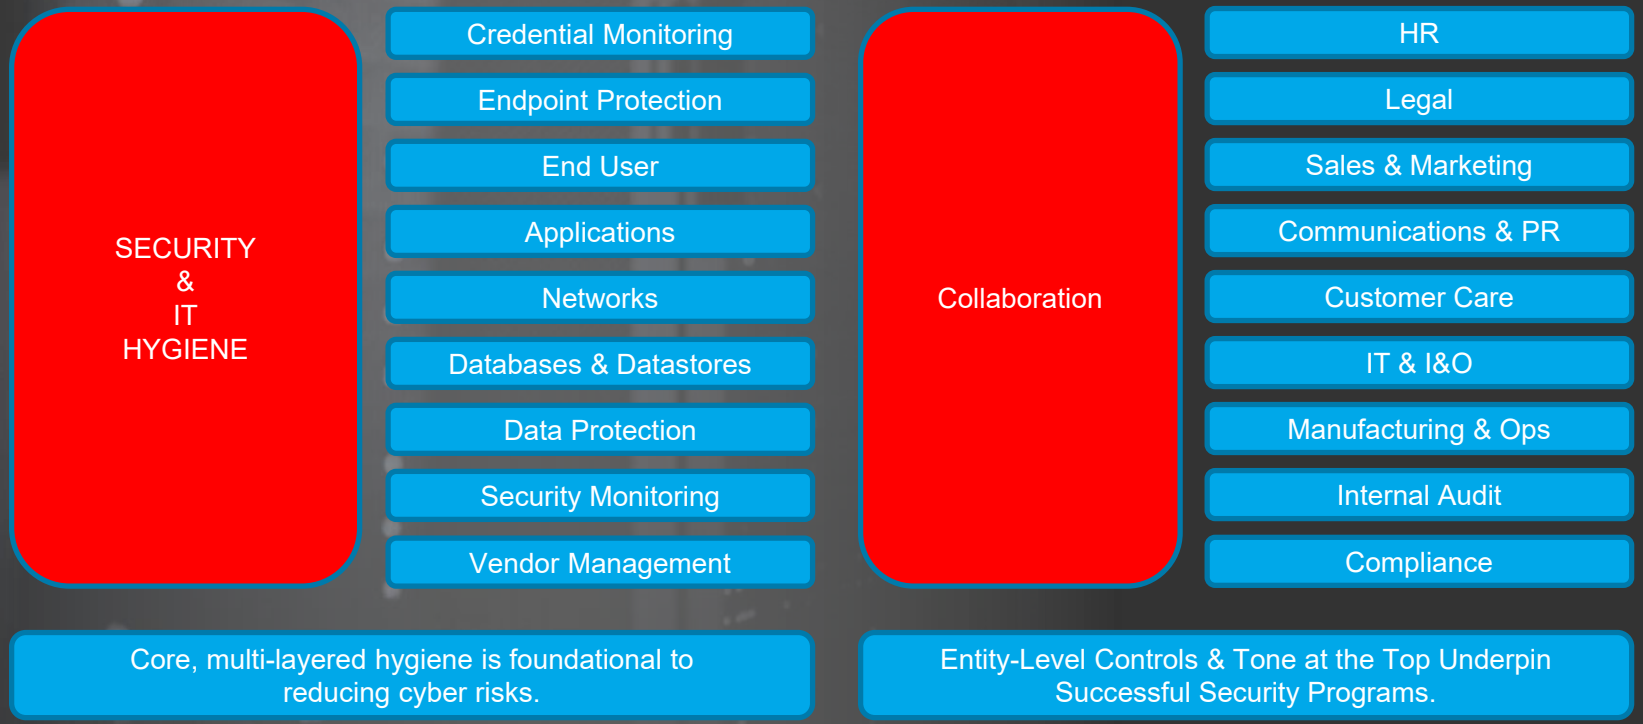

But good security is achievable…

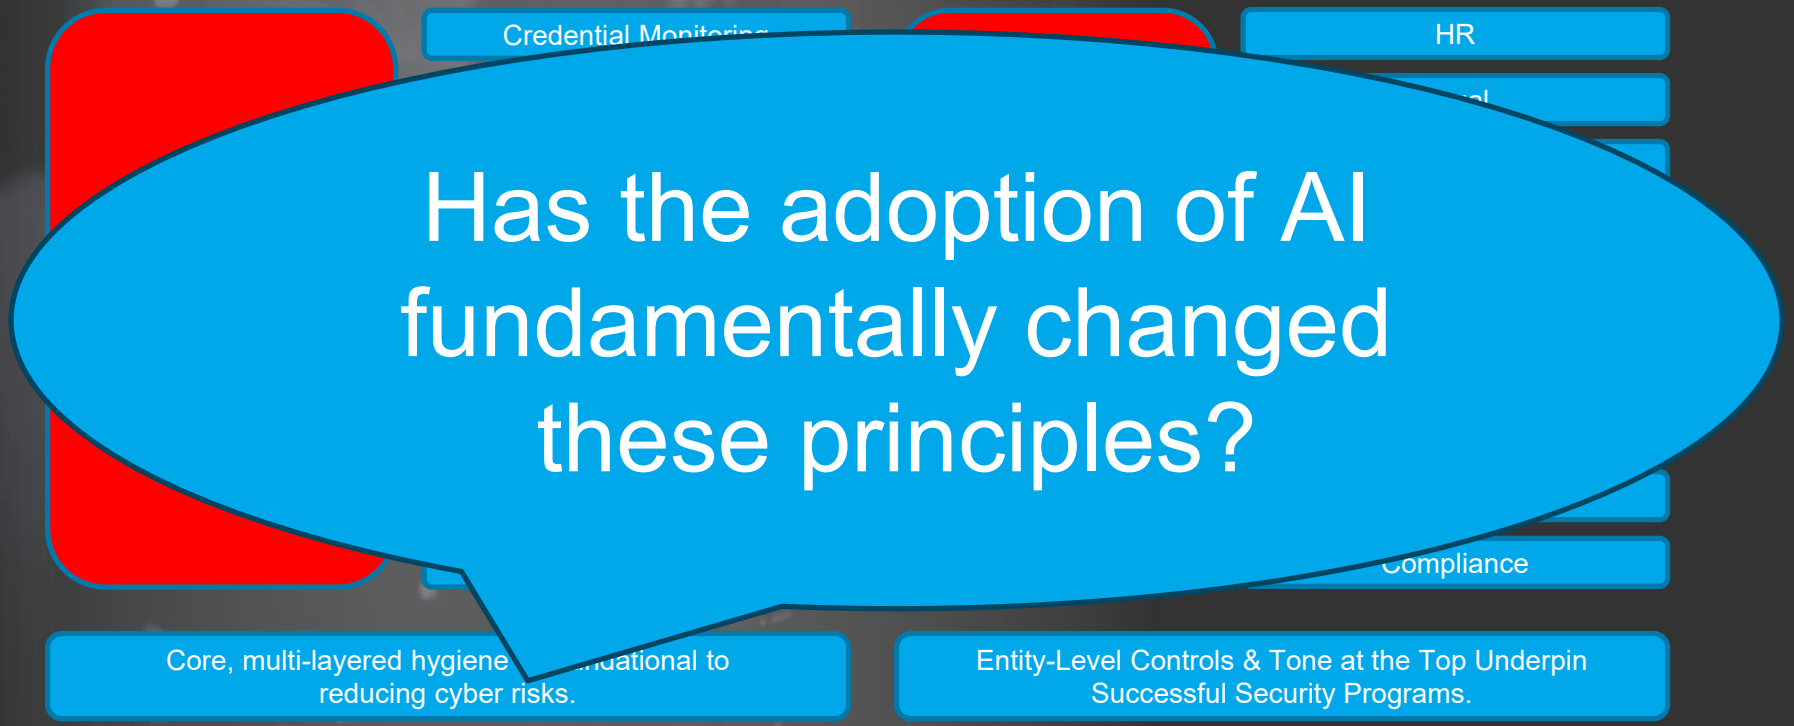

Administrative & Security Program Controls

Reasonable Security – Security Program Capabilities

- Identity Governance
- Asset Management - Fixed
- Asset Management - Intangible
- Vendor Management
- Vulnerability Management
- Policies & Procedures
- Incident Response
- IT Service Management

What is ‘reasonable’ programmatically?

- Quarterly credential & entitlement reviews?
- Least privilege & need-to-know principles?
- On-boarding, off-boarding, and role-change checklists & validations?
- Procedures to find orphaned accounts and clear-text passwords in scripts, code, etc.?
- Password complexity, aging & non-reuse?
- Authentication logging?
- Inventories of service & system accounts?
- Non-repudiation?
- Single Sign-on?

Reasonable Security – Security Program Capabilities

- Identity Governance
- Asset Management - Fixed
- Asset Management - Intangible
- Vendor Management
- Vulnerability Management
- Policies & Procedures
- Incident Response
- IT Service Management

What is ‘reasonable’ programmatically?

- Assets are inventoried & tracked throughout their lifecycle?
- Configuration items (CIs) and other asset detail are managed in a CMDB?
- Assets are tied to defined owners for accountability?
- Assets are configured and managed per policy including hardening guidelines?

Reasonable Security – Security Program Capabilities

- Identity Governance
- Asset Management - Fixed
- Asset Management - Intangible
- Vendor Management
- Vulnerability Management
- Policies & Procedures
- Incident Response
- IT Service Management

What is ‘reasonable’ programmatically?

- Data assets are classified & labeled per policy?
- Security protections for intangible assets reflect classification?
- Data retention schedules are tied to contractual and regulatory obligations?
- Sensitive data is protected by least privilege, need-to-know, and separation of duties principles?
- Cryptographic services are employed per policy and based on classification?
- Access and changes to data are governed and logged?

Reasonable Security – Security Program Capabilities

- Identity Governance
- Asset Management - Fixed
- Asset Management - Intangible
- Vendor Management
- Vulnerability Management
- Policies & Procedures
- Incident Response
- IT Service Management

What is ‘reasonable’ programmatically?

- Vendors are classified based on risk?
- Material vendors are reviewed for risk factors throughout their lifecycle, not just during on-boarding or implementation?
- Service demarcation is clearly identified?
- Roles & responsibilities are clearly identified?
- Remote network connections are inventoried?
- Contracts and service agreements are inventoried and reviewed by counsel?
- Data that is shared with vendors is governed throughout its lifecycle?

Reasonable Security – Security Program Capabilities

What is 'reasonable' programmatically?

How do our core vendors use AI?
What data has been shared with these vendors?
Have our incident response procedures been updated accordingly?

- by counsel?
- Data that is shared with vendors is governed throughout its lifecycle?

Reasonable Security – Security Program Capabilities

- Identity Governance
- Asset Management - Fixed
- Asset Management - Intangible
- Vendor Management
- Vulnerability Management
- Policies & Procedures
- Incident Response
- IT Service Management

What is ‘reasonable’ programmatically?

- Assets are scanned for vulnerabilities with an agreed-to cadence?
- Firmware and software components are scanned?
- SLAs are established for the remediation of identified vulnerabilities?
- Patches and updates are implemented based on risk?
- Assets are categorized based on their enterprise value or inherent risk?
- Vulnerability scoring is contextualized to the operating environment of the asset or service?

Reasonable Security – Security Program Capabilities

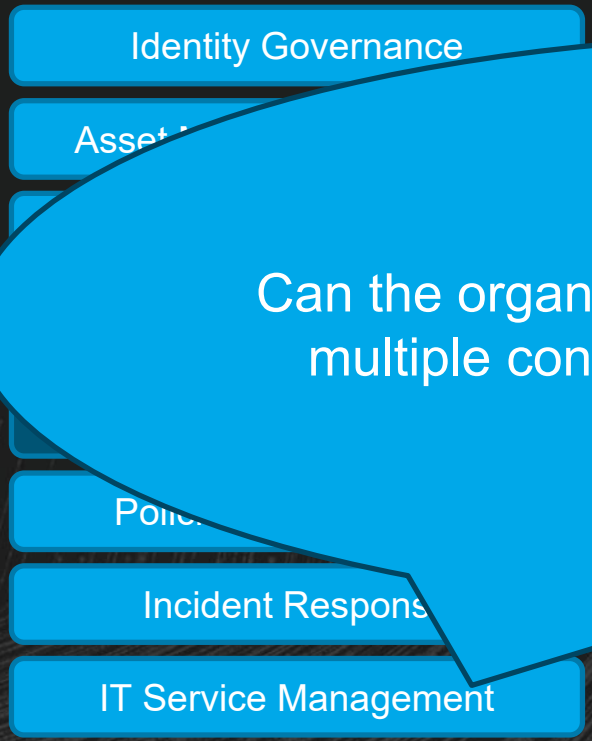

What is ‘reasonable’ programmatically?

with an

Can the organization patch at scale and address multiple concurrent zero-day vulnerabilities?

on

enterprise

Vulnerability scoring is contextualized to the operating environment of the asset or service?

Reasonable Security – Security Program Capabilities

- Identity Governance
- Asset Management - Fixed
- Asset Management - Intangible
- Vendor Management
- Vulnerability Management
- Policies & Procedures
- Incident Response
- IT Service Management

What is 'reasonable' programmatically?

- Policies & procedures are managerially reviewed and approved?
- Policies & procedures are tied to control standards or frameworks?
- Policies & procedures are maintained current with clearly defined next scheduled review dates?
- Training on policies & procedures is maintained current?

Reasonable Security – Security Program Capabilities

- Identity Governance
- Asset Management - Fixed
- Asset Management - Intangible
- Vendor Management
- Vulnerability Management
- Policies & Procedures
- Incident Response
- IT Service Management

What is 'reasonable' programmatically?

- The incident response plan is updated annually?
- Runbooks for common scenarios are drafted:
  - Ransomware
  - Lost or stolen device
  - Credential compromise
  - Malware
  - DDoS
  - Compromised Vendor
  - Data breach
- Tabletop exercises are conducted minimally each year?
- Retained counsel & forensic support?
- Validation of organizational priorities?

Reasonable Security – Security Program Capabilities

- Identity Governance
- Asset Management - Fixed
- Asset Management - Intangible
- Vendor Management
- Vulnerability Management
- Policies & Procedures
- Incident Response
- IT Service Management

What is ‘reasonable’ programmatically?

- Change requests are logged, triaged for risk, and managed through their lifecycle?
- Problems are managed through their lifecycle including identification of root causes?
- Configuration & asset management status is centrally managed?
- Capacity for services is managed consistent with business requirements and obligations?
- Risk and dependency analyses are performed annually?

Technical Controls by Layer

Reasonable Security – Technical Controls by IT Service Layer

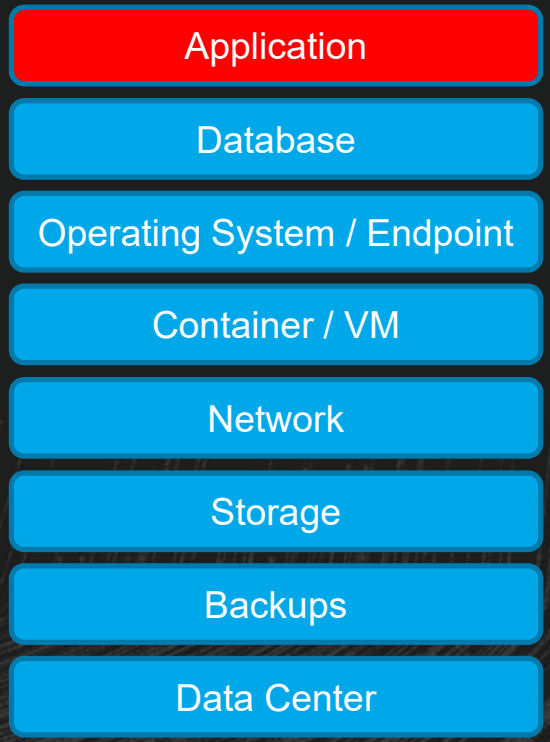

What is ‘reasonable’ at each layer?

- Penetration testing?
- Code reviews?
  - Dynamic Application Security Testing (DAST)
  - Static Application Security Testing (SAST)
  - Runtime Application Security Protection (RASP)
- API Security?
- WAF?
- Secrets management?
- Code signing?

Reasonable Security – Technical Controls by IT Service Layer

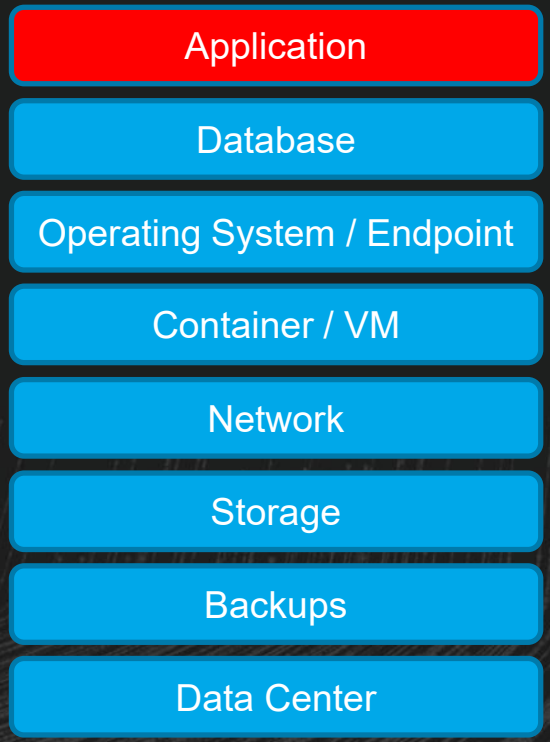

What is ‘reasonable’ at each layer?

- Software Composition Analysis (SCA)?
- Software Bill of Materials (SBOMs)?
- Input validation?
- Error handling?
- Memory Management?

Reasonable Security – Technical Controls by IT Service Layer

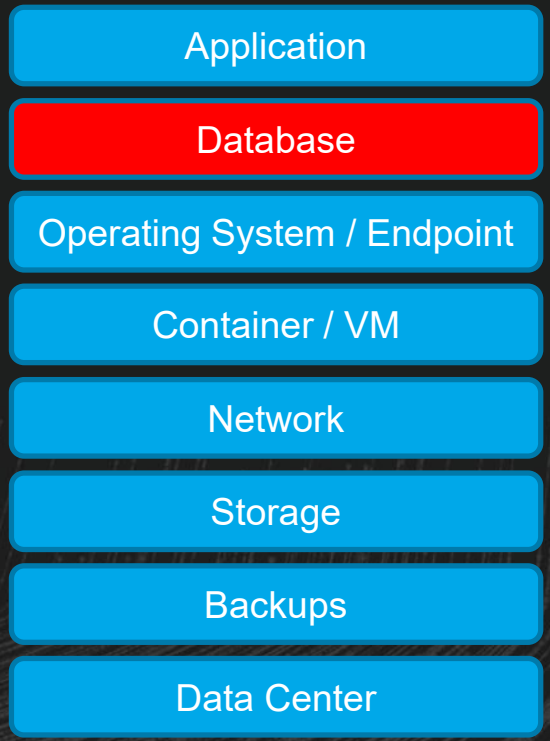

What is ‘reasonable’ at each layer?

- Database Activity Monitoring?
- MFA for database admins?
- Separate credentials for database admins?
- Encryption?
- DLP?
- Hardening & deprecating unnecessary services?
- Logging & SIEM?
- TLS for connections between applications and databases?

Reasonable Security – Technical Controls by IT Service Layer

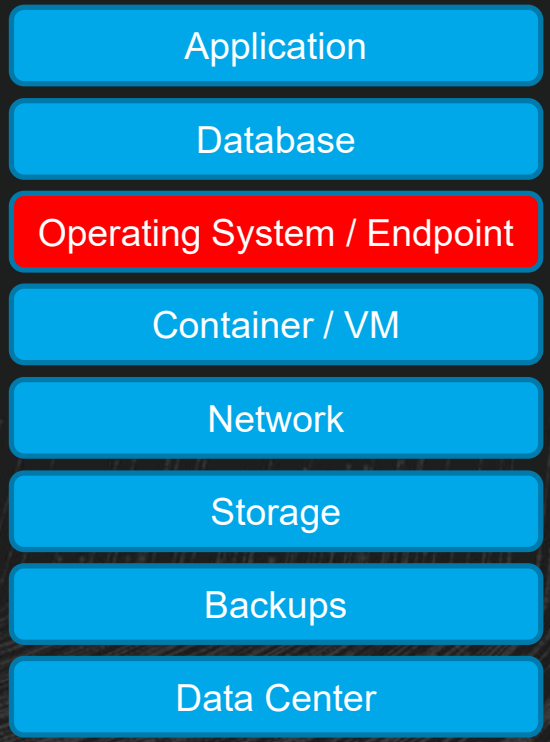

What is ‘reasonable’ at each layer?

- AV, EPP or EDR?
- Encryption (device or file)?
- Mobile device management (MDM)?
- Data loss prevention (DLP)?
- Golden images?
- Host-based IDS?
- VPN?
- Logging & SIEM?
- Deprecated local admin rights?
- Email security?
- Browser & web security?

Reasonable Security – Technical Controls by IT Service Layer

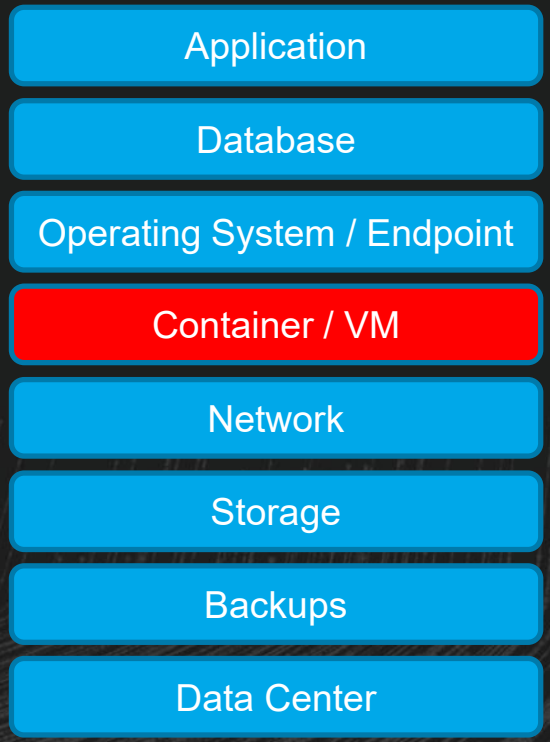

What is ‘reasonable’ at each layer?

- Hardening?
- Root key management?
- Disabling public access?
- Encrypted secrets?
- Implementing networking policies (Kubernetes)?
- Host hardening?

Reasonable Security – Technical Controls by IT Service Layer

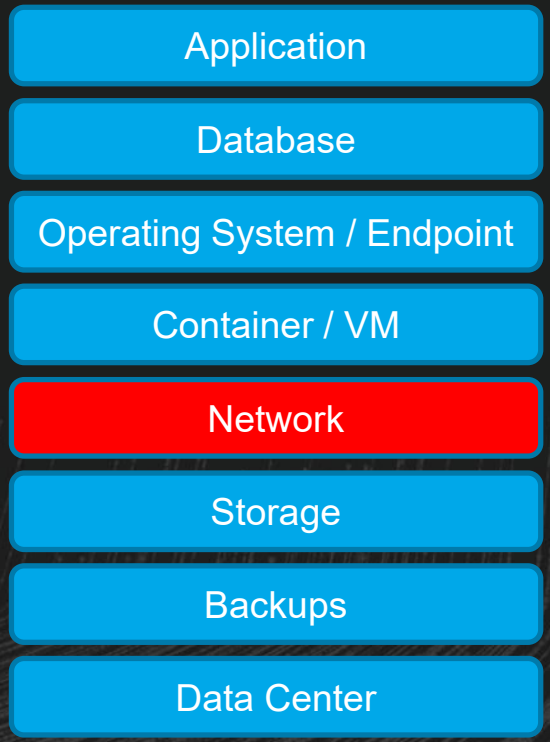

What is 'reasonable' at each layer?

- Segmentation?
- Firewalls (perimeter)?
- Firewalls (internal)?
- Network Traffic Analysis (NTA)?
- Network Access Control (NAC)?
- Anomaly Detection?
- Firmware updates on routers, switches and other network devices?
- Protocol security (e.g., SSL & TLS)?
- Fault tolerance & resiliency of components & ISPs?

Reasonable Security – Technical Controls by IT Service Layer

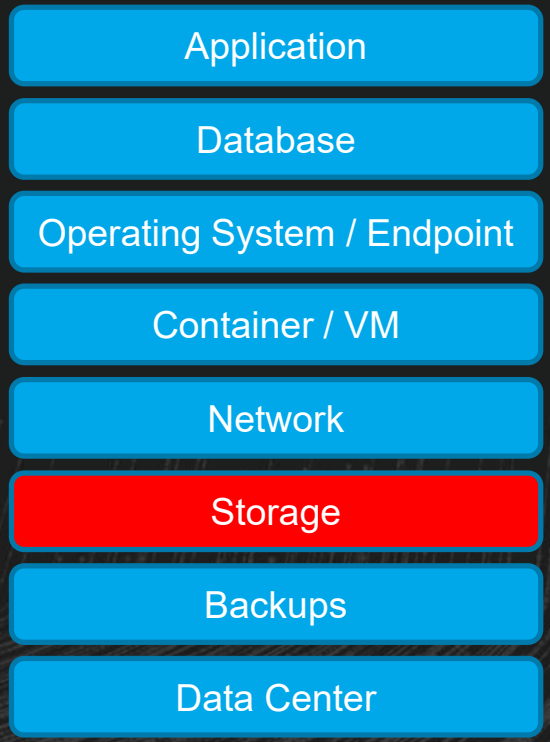

What is ‘reasonable’ at each layer?

- Encryption?
- World-Wide Name, LUN masking & HBA security for SANs?
- Fabric security & zoning for SANs?
- Secure management interfaces (all storage types)?
- Protocol security (NFS, SMB, CIFS, etc.)?
- VLAN & segmentation for iSCSI and NAS?

Reasonable Security – Technical Controls by IT Service Layer

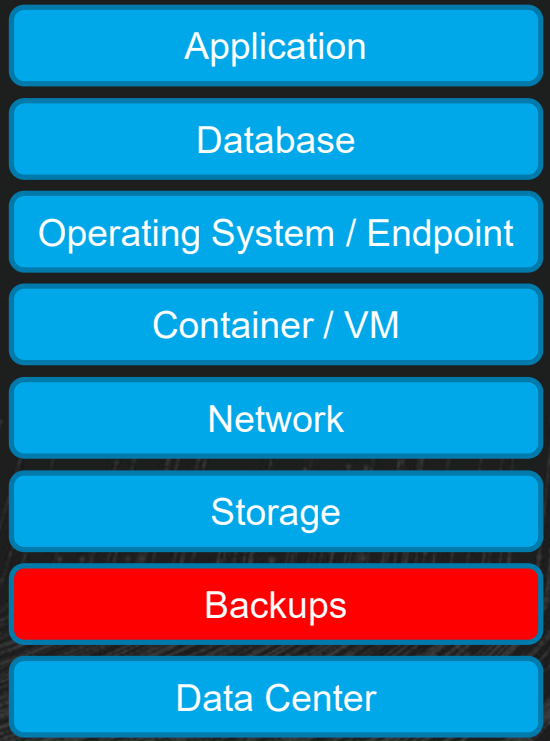

What is ‘reasonable’ at each layer?

- Encryption?
- Media labeling for off-site backups?
- Onsite or out-of-region?
- Physical security over media?
- Backup network protocols?
- Jurisdiction confirmation?

Reasonable Security – Technical Controls by IT Service Layer

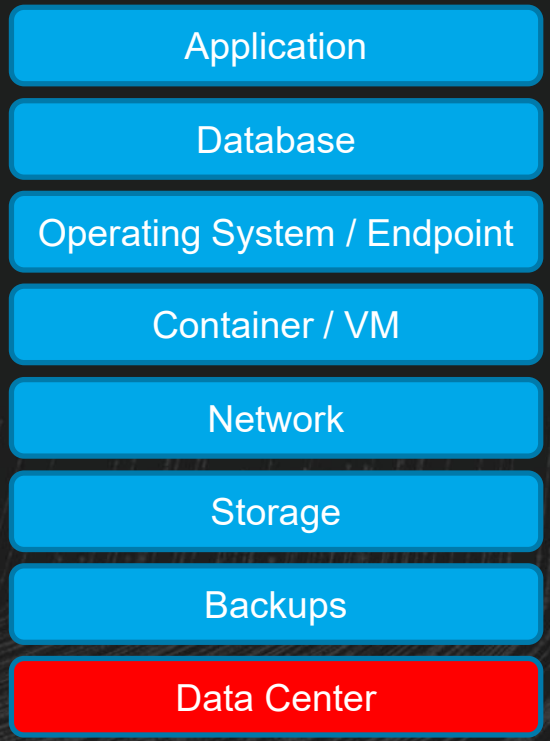

What is ‘reasonable’ at each layer?

- Physical access controls?
- Access & egress monitoring?
- Onsite security?
- Environmental controls?
- Power & UPS?
- Cooling & HVAC?
- Staffing levels & 7x24?
- Number of locations or geographic diversity?


Is AI Security Different?

An Urgent Plea to Start Threat Modeling AI

How Should Our Threat Models Change?

Threat Modeling

STRIDE is an efficient way to analyze the security risks of applications and technology.

- Spoofing
- Tampering
- Repudiation
- Information Disclosure
- Denial of Service
- Elevation of Privileges

Threat Modeling

STRIDE is an efficient way to analyze the security risks of applications and technology.

- Spoofing
- Tampering
- Repudiation
- Information Disclosure
- Denial of Service
- Elevation of Privileges

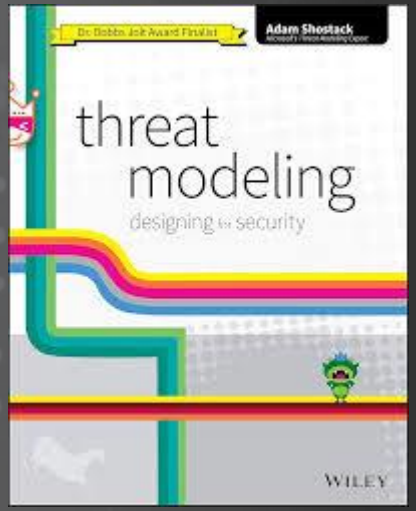

AI Threat Modeling

The Cloud Security Alliance and other organizations have excellent resources to facilitate threat models contextualized to AI-specific risks.

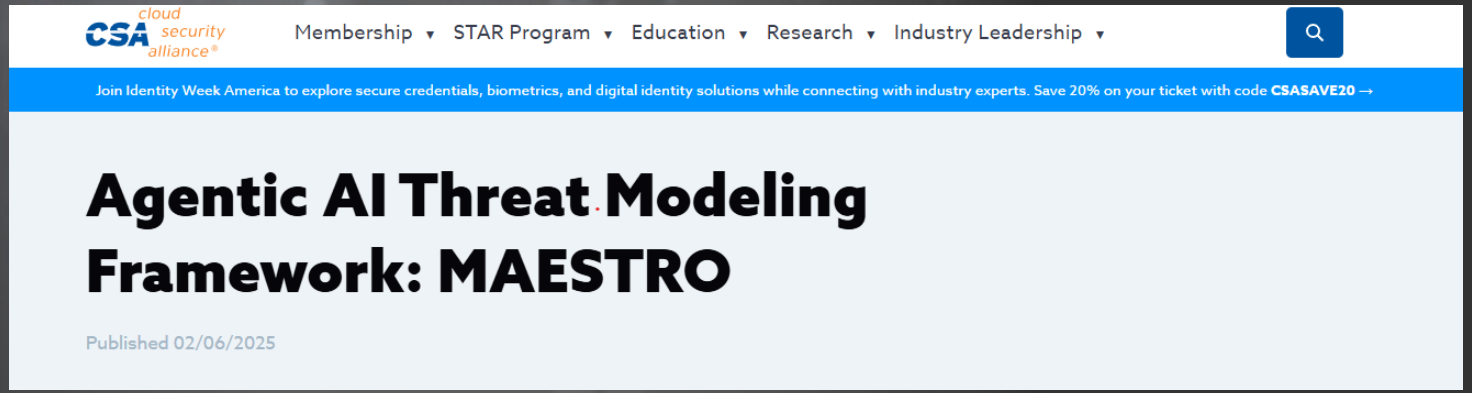

Model Context Protocol (MCP)

MCP is used to coordinate the activities between disparate applications and language models.

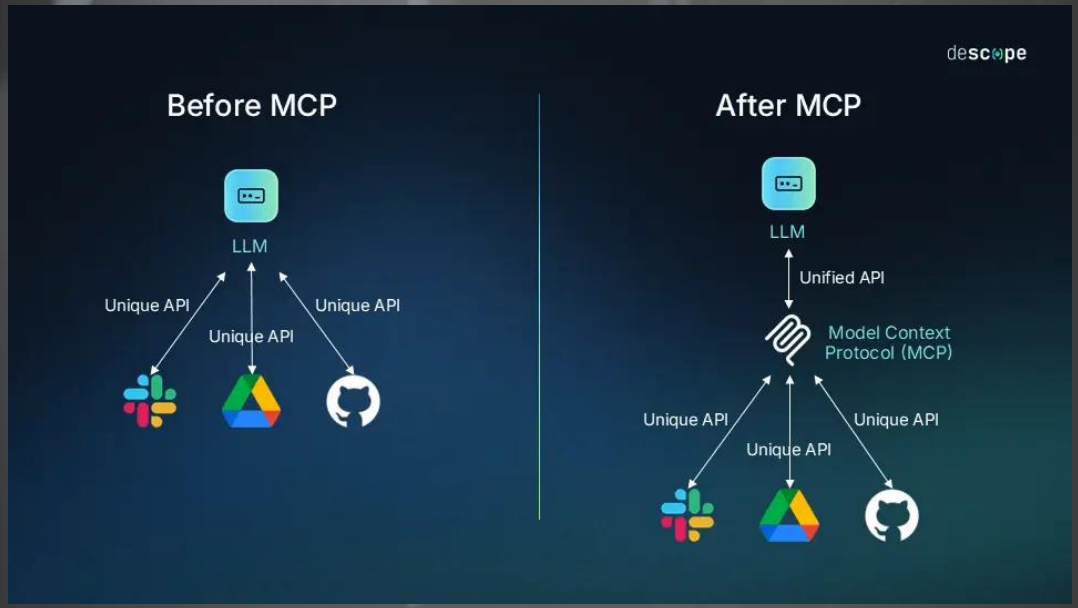

Source:

https://www.descope.com/learn/post/mcp

(No endorsement / just the graphic)

Model Context Protocol (MCP)

MCP is used to coordinate the activities between disparate applications and language models.

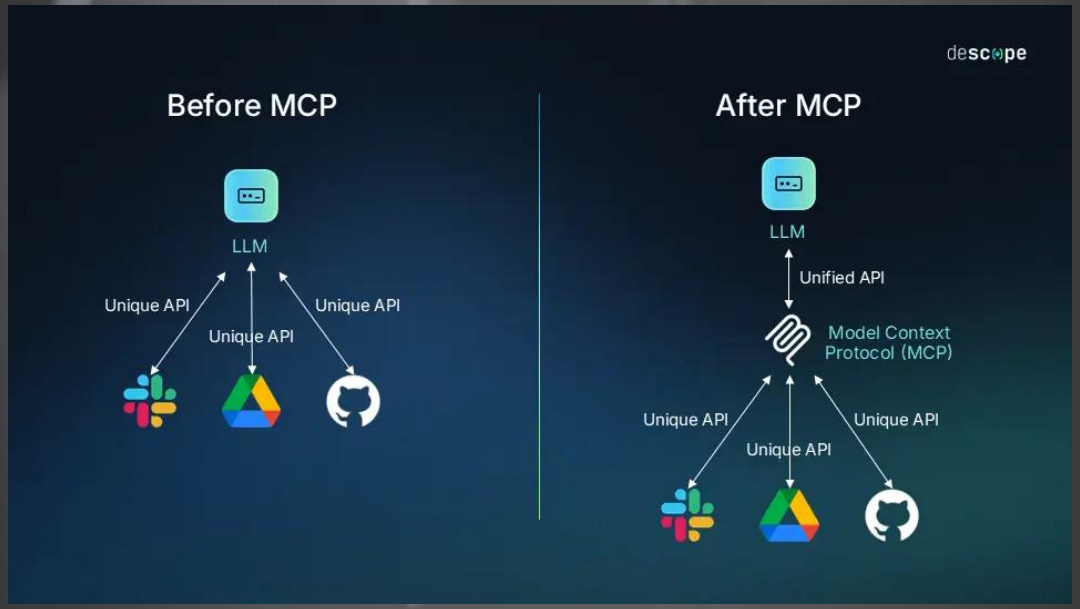

We need to secure the integrations:

- SBOM
- SCA
- DAST
- SAST
- Runtime
- APIs

Agentic AI

Agentic AI is changing how we do work.

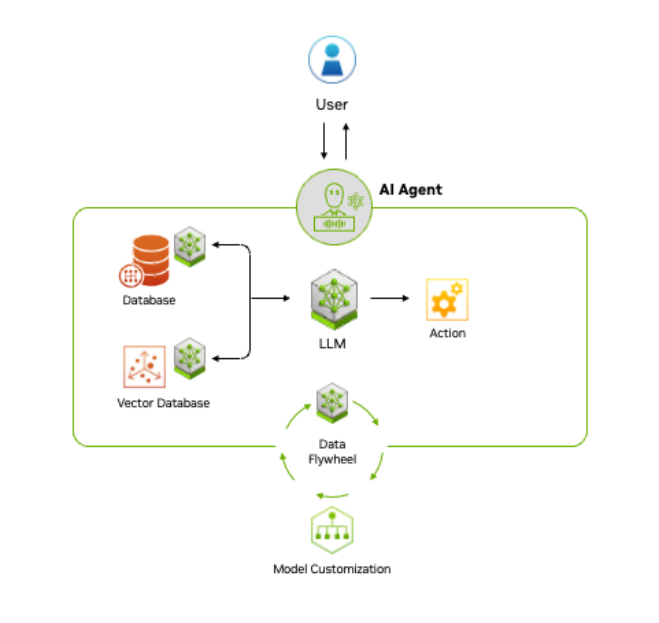

Source:

https://blogs.nvidia.com/blog/what-is-agentic-ai/

Agentic AI

Agentic AI is changing how we do work.

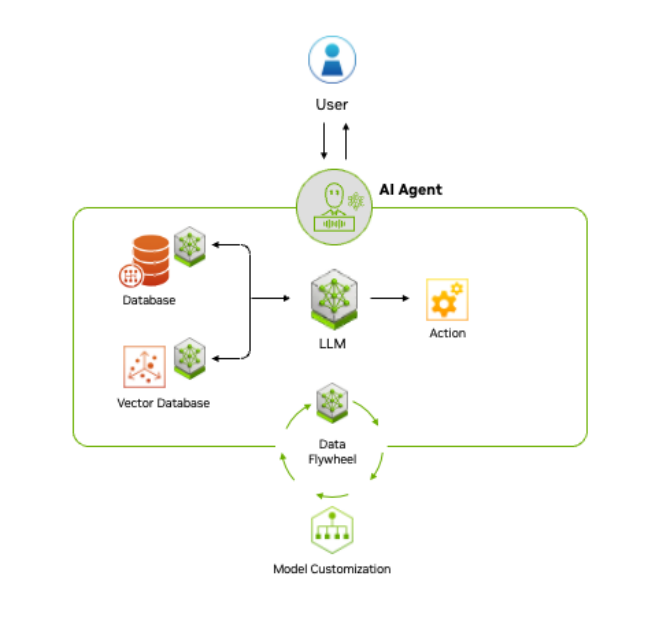

What happens when:

- Agents go rogue?
- Agents are compromised?
- Agents lack context?

- Other risks?

Vibe Coding and 'Citizen Developers'

Today, everyone can create an application. What are the implications to:

- Code quality?
- LLM Authorization?
- Attack surface?
- Privacy?
- Safety?
- Third-Party Dependency Analysis?
- Provenance of Models?

FMC – A Lesson Learned from Gartner

Cybersecurity is tough. Securing AI technologies and services will likely be even tougher given our current track record. Cybersecurity absent good governance and an organizationally-aligned target operating model is 'nigh impossible.' Some key recommendations:

- Ensure roles and responsibilities are clarified
- 'Declare War on Ambiguity'
  - Dependencies
  - Resources
  - Licensing
  - Sourcing Model
  - Competencies
  - KRIs
- Align Security Functions to the Business

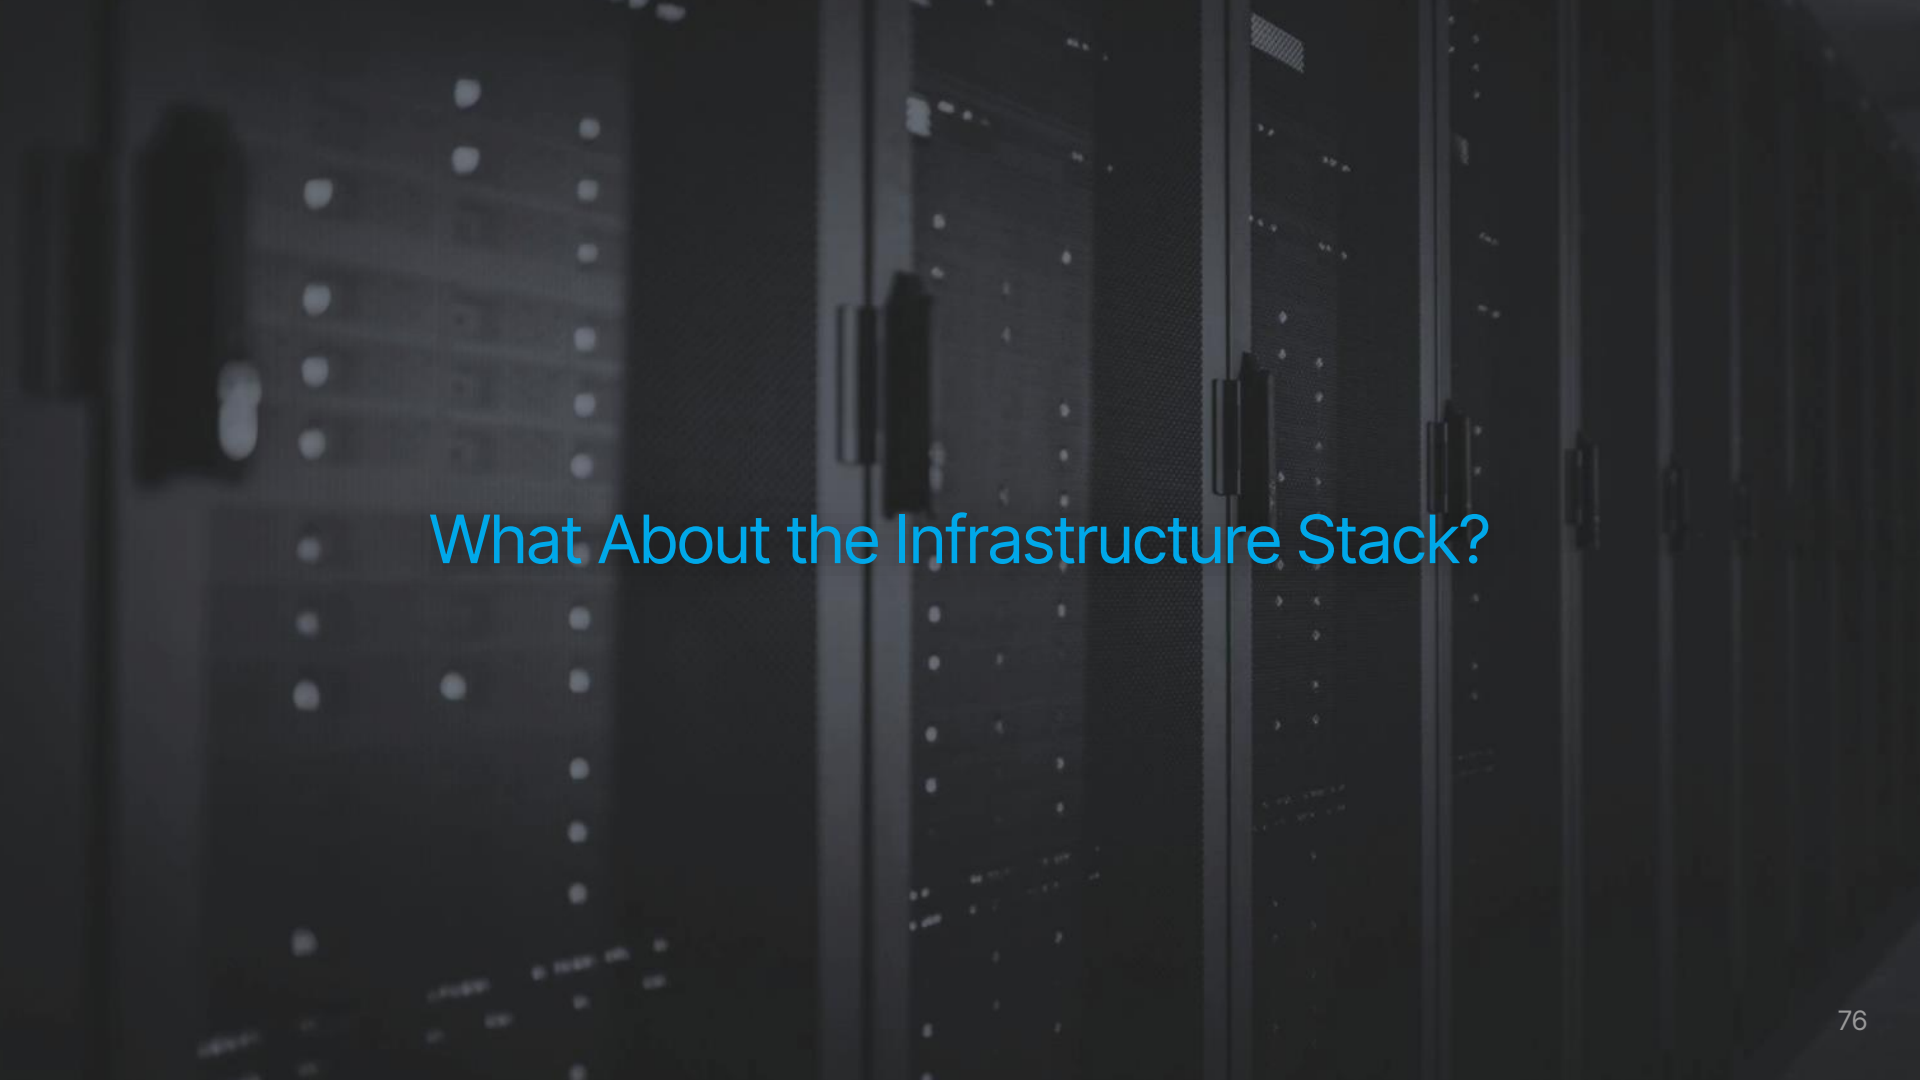

What About the Infrastructure Stack?

Reasonable Security – Technical Controls by IT Service Layer

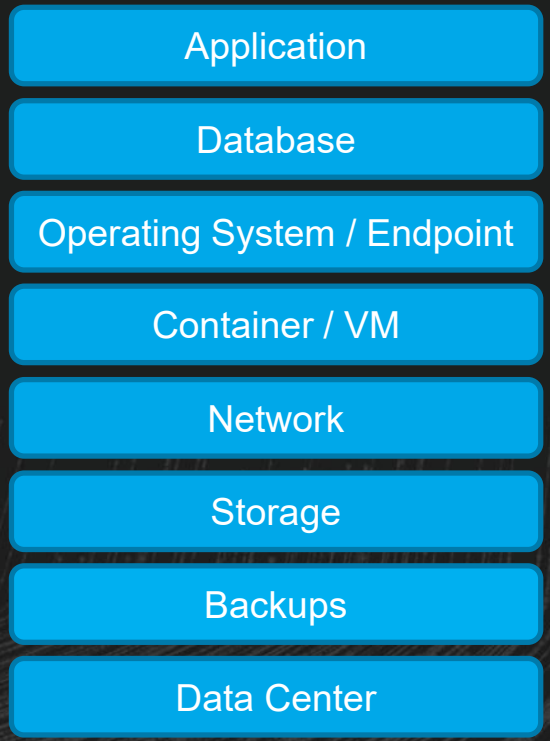

What is ‘reasonable’ at each AI layer?

We need to understand if and where AI applications differ from current applications.

- AI in SaaS?
- AI in Data Governance?
- AI in API Security?
- AI in Microservices Security?
- AI in Code?
- Algorithms?

Some final thoughts on AI security governance…

- Ensure there is an identified and authorized security leader
- Determine the appropriate AI security strategy and target operating model
- Declare war on ambiguity and use key risk indicators (KRIs) and other metrics to evaluate coverage, operational risk reduction, and personnel risk factors
- Select an agreed-to security standard or framework (e.g., NIST AI Risk Management Framework, NIST CSF, ISO, 42001, or ISO 27001)
- Create a security calendar and ensure that security functions are fully captured in the calendar and resourced accordingly
- Conduct a business impact analysis (BIA) to help contextualize security capabilities to enterprise value and capture key dependencies (technology, staff, locations, vendors, etc.)
- Use a risk register to track identified strategic risks and link technical and security risks to enterprise risk factors
- Validate, validate, validate!

Thank You!

Matt Stamper, CEO | CISO Advisor, Executive Advisors Group, LLC
MPIA, MS, CISA, CISM, CRISC, CDPSE, AAIA, CIPP/US, ITIL
matt@executiveadvisorsgroup.com