NYC Cyber Command

Request for Information (RFI) for:

Zero Trust Security Management

Moving Beyond the Perimeter

TABLE OF CONTENTS

1. RFI Purpose and Objectives

1.A. RFI Purpose

Protecting the City of New York's information infrastructure is vital to the proper functioning of the City and the ability for agencies and personnel to serve the residents, visitors, and businesses of the City of New York. As an enterprise, NYC requires a security model that works with the complexity of our highly diverse and federated architecture, accounts for an increasingly mobile workforce, and protects people, devices, applications and data wherever they are located. The traditional "castle and moat" approach will not create or sustain the future of our cyber resiliency, and we want your ideas about how Zero Trust Architecture solutions can play a role in that future.

The New York City Cyber Command (NYC3) is exploring how a Zero Trust Architecture can be implemented in a manner that is tailored to the City's unique infrastructure, and ultimately improve the security posture of the City of New York. Zero Trust components will confront common public sector challenges:

- Technical debt that may preclude micro-perimeters and APIs
- Legacy systems with potentially non-compatible applications
- Dispersed oversight among our hundreds of agencies and offices
- Significant PII and similarly sensitive data ownership
- Contractor involvement across the City that may involve external Identity and Access Management (IAM) systems

Our aim is to address the City's fragmented identity and resources through a staged deployment: unifying IAM for applications and servers, followed by implementing context-based access, and finally, instituting an adaptive Zero Trust Architecture that relies on risk-based access and continuous authentication.

We want to hear your ideas for how Zero Trust can work, and what kind of pilot approaches make sense to you. In this RFI, we are seeking your vision of a strategy that is manageable, cost-effective, and handles complexity at the scale of New York City. Let us know how you would sequence, stage, and implement that vision; what key barriers you envision; and what you can (and cannot!) deliver.

This RFI is a starting point for conversations with vendors before we determine our implementation plan. Respondents can answer any or all questions we have included in this RFI. Organizations may partner and submit joint responses. Participation in this RFI is not required for any future RFP and will not help or hurt your ability to partner with us in the future. This RFI is not a solicitation for services or products. There will be no contract award(s) resulting from any submission in response to this RFI.

Please provide a response in no more than 100 typed pages, including appendices. You can refer to the following Section 2, "Current State," for more information about NYC's environment and ongoing cybersecurity initiatives. By submitting a response to this RFI, the Respondent authorizes the City to access and utilize all information provided without limitation or condition.

1.B. About NYC3 and NYC Cybersecurity

In accordance with the City Charter, NYC3 serves as the body accountable for the City's centralized cyber defense to protect City resources, employees, and the public from cyber threats. NYC3 is tasked with ensuring Citywide agencies are in compliance with information security policies and standards. Our Office also deploys technical and administrative controls to mitigate cyber threats.

NYC3's mission is to lead and execute an innovative, intelligence-driven, risk-informed cyber defense and response strategy, which enables the City government to properly function and provide services to New Yorkers. Our vision is that New York City is the most cyber-resilient city in the world so that the services which New Yorkers rely on are available when they need them.

OneNYC 2050 is New York City's long-term strategic plan. Included in OneNYC 2050 are efforts to improve the City's digital infrastructure, and to build and cultivate an innovative cybersecurity ecosystem alongside City partners. Other initiatives include improving the City's data infrastructure to enable greater data integration and agency collaboration.

NYC3 is not new to Zero Trust architecture. Our Office has built a highly secure responder environment based on the principles of Zero Trust. NYC3's environment is designed to ensure the security and reliability of critical systems and services through a Zero Trust environment and supporting architecture. This has helped NYC3 ensure a continuity of operations, including a zero-downtime switch to 100% remote work during the COVID-19 pandemic.

While recognizing the benefits, the citywide implementation of a Zero Trust Architecture presents a wide variety of challenges due to the federated digital environment. As a result, NYC3 believes engaging the public through this RFI can help determine how Zero Trust can be rolled out across the remaining City Agencies in an efficient, scalable manner.

2. Current State

NYC3 is driving initiatives across the City's Identity and Access Management (IAM), endpoint, and application environments to ensure it is meeting the latest best practices. Below, we describe some of these efforts to help you structure your response on how a Zero Trust Architecture could enable us to meet our cyber resiliency goals.

2.A. Identity Environment

Overview

IAM is an integrated system of capabilities that enhances information security, reduces administrative overhead, and improves timelines of providing and reporting access. NYC's IAM program is known as NYC.ID, the brand name for the City's centralized Lightweight Directory Access Protocol (LDAP) service. It is the repository for account information, the identity provider (IDP) for the City, and is a Department of Information Technology and Telecommunications (DoITT) service offering. There are other IDPs throughout the City.

The goal of the internal City workforce IAM program is to ensure all human accounts will be initiated, transitioned, and terminated by business process drivers. Employees' identity and authorization are vetted and verified to ensure that accounts are accurate and updated, and employees have authorization to the resources needed to perform their duties. We aim to ensure that access to critical data and functions is protected by strong authentication, auditing, and recording, but know that opportunities still exist to improve the standardization of New York City's identity environment.

A number of City workforce user types exist within the City's identity environment, such as:

- City contractor staff,
- Business partner users,
- Agency-trusted users,
- Non-human users,
- Active employees with multi-agency assignments,
- Active City agency employees,
- Inactive City agency employees still on payroll,
- and departed City agency employees.

Agencies have deployed IAM mechanisms including Microsoft Active Directory, and cloud-based tools such as Microsoft Azure Active Directory, Google Workspace (G Suite), and others. The City is pursuing options to continue to improve account provisioning and identification efforts.

Current Policies Influencing IAM:

DoITT's ***External Identity Management and Password Policy*** (published Fall 2014) requires public facing applications use DoITT-provided identity management services from NYC.ID and NYC.gov to automate user registration and authentication. Each agency is responsible for the management of its user identities, including identity validation and registration, authentication, authorization, provisioning and deprovisioning of identities. Management approval is required before a user is authorized to use any City computing services. Employees, as well as consultants who are working with the City under a contractual agreement, may have access to City computing resources if they have a nondisclosure agreement and the sponsoring agency approves their access.

Under ***DoITT's*** ***Identity Management Security Policy*** (published Fall 2014)***,*** users are authenticated at a level corresponding to the data classification of the information they need to access to fulfill their work requirements. Under this policy, access permissions and data classification are defined, and standards are created to allow agencies to follow recommended guidelines to authenticate users. Lastly, user accounts will be created and de-provisioned in a timely manner and inactive user accounts will be de-provisioned.

DoITT's ***Password Policy*** (published Fall 2014) defines user, administrative, and service accounts, with administrative accounts being provided to individuals with the need to carry an elevated degree of privileges such as managing systems, user accounts, and password resets.

2.B. Endpoint Environment

Overview

The endpoint solutions used by the City of New York currently aim to protect agencies from threats to endpoints by prompting responsive control to alerts and decreasing the endpoint footprint to the user and administrator by downsizing identities (in a federated environment) and third-party requirements.

The Current State of Endpoint Protection Tools

NYC3 is transitioning the City towards a centrally managed security approach. The City primarily uses endpoint detection and response technology on its endpoints, though anti-virus, whitelisting, host-based firewalls, and other methods and technologies are also used. There is some variation in how agencies have configured their endpoint security, and which modules are active. The Endpoint team at NYC3 has created endpoint security requirements customized to the City's unique, multi-tenant and maturity-variant environment. The NYC3 Endpoint Security team uses the term "**Target State**" to identify the requirements for a secure endpoint and design technology around the following **eight domains**:

Target State		Description
1	**Asset Management**	The asset management category covers overseeing the tracking, maintenance, and classification of the City of New York's various endpoints. As these endpoints process, transmit, and retain sensitive user information, and run systems critical to maintaining the City's operations, it is vital to maintain an accurate listing of these endpoints. This includes governing what systems are installed on the endpoint, categorizing the endpoint based on its function, and maintaining and tracking it throughout its lifecycle.
2	**Configuration Standards**	Baseline configurations for endpoints provide a secure template that is used to compare possible future modifications. This requirement focuses on ensuring a baseline image standard exists and is maintained with hardening standards and essential functions. All supported operating systems must have a current image that is used when provisioning new endpoints, as this streamlines the process and provides new endpoints with the necessary protective security features.
3	**Authentication and Authorization**	These requirements cover the process of users gaining access using non-domain credentials to New York City endpoints. Domain access is not covered, as domain controls and requirements are handled outside the realm of the endpoint. These requirements cover local accounts, service accounts, shared accounts, and password policies and requirements.
4	**Connected Device Protection** (*Removable Media*)	This requirement covers removable media interaction, such as a thumb drive.

		All interactions must be logged and monitored, and the capacity of use can differ based on risk and functionality. However, auto-execute must always be forbidden.
5	**Encryption**	Drive encryption provides a line of defense if an endpoint's hard drive ends up in the possession of an unauthorized user. These requirements focus on encrypting the City's endpoint drives and provisioning endpoints with encrypted drives. Credentials and processes for initial encryption, as well as key storage, must be entirely independent of the endpoint.
6	**Threat Prevention**	Host-based firewall configuration, such as denying all inbound traffic, is intended to include a justification for variances from the baseline, and clear processes and procedures for adding, removing, and maintaining firewall rules. Host intrusion prevention systems (HIPS) protect critical systems from viruses and malware. These requirements are specific to managing and overseeing HIPS, with rules designed to be as specific as possible. Processes and procedures covering addition, removal, and modifications must be documented.
7	**Vulnerability Management**	These requirements encompass all monitoring and logging aspects of endpoints. Logging must be enabled, functioning, and ultimately be sent to a centralized repository. Clocks on the endpoint must be synced to the same source, to ensure the integrity of timestamps. These requirements cover management and mitigation of endpoint vulnerabilities. Endpoints must be compliant with patches and be within the support structure of vendor security patch cycles. Operational procedures must be in place to maintain, and continually patch endpoints and in response to a critical vulnerability discovery.
8	**Detection and Response**	The endpoint investigation operations requirements cover the generation, allocation and assessment of security alerts. Process and procedures must be in place that define how to utilize the data, and who will be executing the response.

Endpoint Protection Governance

To stand up to this target state and oversee endpoints, a governance operating model and RACI (responsibility assignment matrix) has been designed. These governance functions will help enforce compliance and address the shared responsibility of endpoints between Cyber Command and external agencies.

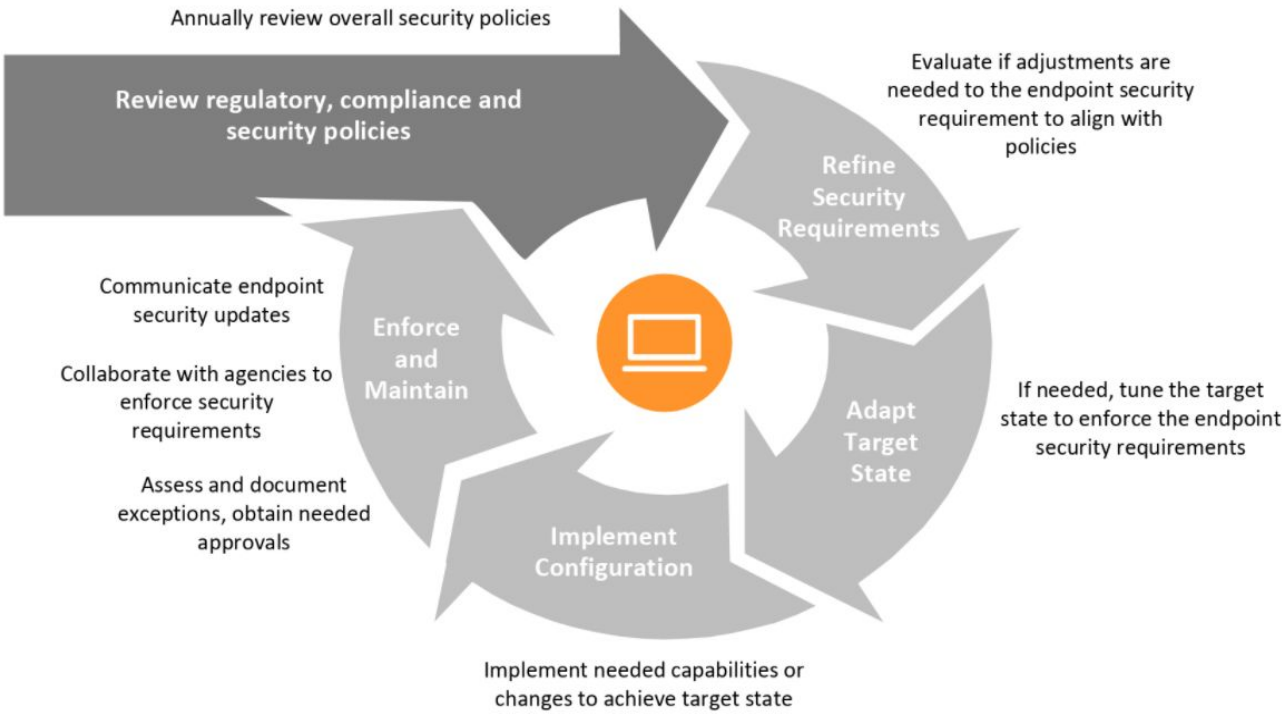

2.C. Application Environment

Overview

The application environment for NYC agencies can be categorized into **three stages of deployment**: the production environment, the staging environment, and the development environment.

Stage of Deployment		Description
1	**Development Environment**	Development is the first stage of the deployment cycle, where engineers and application developers focus on creating and testing the environment and code to make sure the application works correctly. It is the most basic environment for testing, primarily for debugging.
2	**Staging Environment**	Staging is the second stage of development and deployment of the application. This environment is configured to run as close to the actual production environment as possible (including firewalls, scaling, data, etc.). The purpose of this environment is to get a working version of the application approved by the application manager and stakeholders before being launched into production.
3	**Production Environment**	Production is the last stage of development, when the application is live internally or externally for third parties. This environment is the most sensitive to disruptions. Applications must undergo extensive User Acceptance Testing (UAT) and tested in an

		integrated development environment (IDE) before they can be deployed into a production environment.

In order to become authorized as an application on the City's network, applications must perform a unit, system, security and user acceptance test case against the staging environment. The application assessment is designed to identify any unanticipated interactions with existing systems prior to moving to the production environment. Reassessment frequency is based on the criticality of the application.

- High Criticality - assessed every six months
- Medium Criticality - assessed every twelve months
- Low Criticality - assessed every eighteen months

Software applications and services at NYC agencies are enrolled in Software Security as Assurance (SSA). The SSA process ensures that software is designed, implemented and deployed to operate at a level of security that mitigates depreciation of the confidentiality, integrity and availability of the system or its data. SSA translates security requirements into tasks that are implemented throughout the different phases of the Software Development Life Cycle (SDLC) phases, so that City applications are secure by design.

Commercial Off the Shelf (COTS) applications are not developed specifically for the City of New York. As a result, COTS applications may not be fully customizable to the City's standards. The City uses a combination of vetting, security assessment and expectations in order to ensure that COTS applications meet its threshold security requirements.

- COTS Vendor Vetting - defined security requirements that inform the vendor selection process
- COTS Security Assessment - prior to release, applications that store/pass City data are validated by undergoing a security assessment (defined by NYC Cyber Command)
- COTS Exceptions - applications that do not comply with standard vendor vetting and security assessment requirements must be approved by the Citywide CISO

3. Future Vision

Please describe your vision of Zero Trust implementation across the City of New York's identity, endpoint, and application environments. The future vision should also include reference to relevant technologies, resources, and processes for Zero Trust implementation on the City's network. Responses should be no more than 100 pages including appendices, and should be as specific as possible about the scale, scope, and constraints of your proposals. We are particularly interested in pilot ideas or designs.

While we welcome creativity in your responses, please keep in mind the topics of interest outlined below. In addition, Section 5 of this RFI has been provided as an appendix of sample questions that may relate to your proposal's technical model, business model, and user experience.

Architecture and General System: NYC3 is eager to hear how you would envision the general architecture of your model working within a federated system with significant legacy applications. Please specify what assumptions are included in your model regarding City assets, network protocols

and conditions, and endpoint architecture. We are also interested in hearing how your model would handle encryption and security, what third parties or partners you work with, and what the licensing and use constraints are within your model.

Risk-Based Authentication/Authorization and Identity Management: The City currently relies on several identity management systems. Please describe what kind of authentication standards you would rely on in your Zero Trust model, and what kind of analytics you would perform on identities and users. We are interested in hearing how you would handle a multiple identity environment, and how your trust or policy engine/broker would work.

Risk-Based Endpoint Management: Please describe how you would handle endpoint management. Be sure to include how legacy devices fit into your management solution. To the extent practicable, please describe any secondary products or vendors you work with, and how you would handle the City's diverse endpoint agents and devices in your Zero Trust model.

Monitoring and Continuous Improvement: The City is looking to explore ways of experimenting or piloting models for Zero Trust. We welcome descriptions of how you would roll out your model, and what kind of inspections or monitoring would be involved in your proposal. It is helpful to hear how vulnerabilities would be addressed and communicated, and how you would manage the user experience during roll-out.

Failure Management and Correction: We are interested in how you would manage business disruption, latency, and failure within the context of your proposal. Please describe what your approach would be to investigate, remediate, and otherwise communicate and handle malicious activity and compromised points of failure within your proposed model.

4. RFI Response Instructions

NYC3 is seeking responses on possible Zero Trust models that include:

1) **A Respondent Profile**
2) **Proposal Summary.** Please refer to Section 3 of this RFI for topics of interest, and Section 5 for an appendix of accompanying questions.
3) **Relevant Business Parameters**
4) **Supporting Files**

Responses should be submitted online via email to zerotrust@cyber.nyc.gov. While you are welcome to respond to all or only some of the questions mentioned, all respondents should include their Respondent Profile with the requested details and their relevant experience. Please review the NIST Special Publication 800-207 on Zero Trust Architecture to help inform your response to NYC3's RFI.

~~All submissions must be received no later than 01/31/2021, 11:59PM (EST)~~. The due date for RFI responses has been extended. All submissions must be received no later than **11:59pm (EST) on Sunday, February 28, 2021**.

Respondents may submit questions and requests for additional information concerning this RFI. Questions deemed appropriate will be answered publicly and made available on NYC3's website (https://www1.nyc.gov/site/cyber/collaboration/collaboration.page). All questions should be submitted to zerotrust@cyber.nyc.gov no later than 01/31/2021, 11:59PM (EST).

NYC3 reserves the right to conduct subsequent information gathering or other activities with all or some respondents in keeping with relevant laws, policies, and regulations, at its sole discretion.

Confidential and Proprietary Information

NYC3 will endeavor to protect from disclosure any confidential and/or proprietary information the Respondent submits related to this RFI in accordance with applicable law. Respondent must identify those portions of the RFI response it deems to be confidential, proprietary information, or trade secrets.

Respondents should be aware that NYC3 may be required, pursuant to the New York State Freedom of Information Law ("FOIL") (New York Public Officers Law Section 87 et seq.), to disclose to the public a written response to the RFI or portion thereof.

If such disclosure is requested by a third party, NYC3 will notify the Respondent as practicable of any deadline to respond. Consistent with the requirements of FOIL, NYC will make the final determination of whether such information may be withheld from disclosure. If NYC3 determines that information may not be withheld, we will attempt to provide the Respondent with timely notice of intent to disclose, so that the Respondent may invoke any rights or remedies to prevent disclosure to which it believes it may be entitled to under the law.

Respondents expressly acknowledge and agree that neither NYC3 nor the City of New York will have any obligation or liability to any Respondent in the event of disclosure of materials designated as confidential or proprietary.

Not withsstanding the above, it is the intent of NYC3 to publish summary information about the responses received, including but not limited to the number of responses, names of respondents, and the context of their response.

---

1) Respondent Profile

Please provide a respondent overview that describes your organization and addresses your organization's information related to your response to this RFI. For joint responses, please provide a profile for each organization involved.

Name:

Title:

Organization Name:

Street Address:

City:

State:

Zip Code:

Country:

Phone Number:

Email Address:

Please provide a short statement describing your organization:

2) Proposal Summary

Please provide a statement of your proposal related to the scope, resources, and other parameters you envision for implementing a Zero Trust Architecture model within the City of New York. We are particularly excited about ideas and recommendations that focus on pilot projects that NYC3 could undertake.

We encourage you to provide as much detail as possible about the areas of interest mentioned in Section 3 of this RFI, and elaborated upon in the questions provided in Section 5.

3) Business parameters involved in your proposal

Please describe how you would imagine your organization working with New York City Cyber Command and the City of New York, and please include specific expectations on how your model would address:

- Network architecture updates
- Finance
- Deployment processes
- Maintenance and operations
- Service delivery and customer support
- Performance monitoring, data collection, and public reporting
- Other not listed above

What City collaboration commitments would be important for your proposal to work? Are there any legal or regulatory constraints that are relevant to your proposal?

4) Supporting files

Please include any supporting files, including diagrams, maps or other content to help demonstrate solutions.

5. Sample Questions on Zero Trust

Please refer to the list of sample questions below for additional context to inform your RFI response. We provide these questions to represent the breadth of topics that are of interest to NYC3 as we examine possible Zero Trust Architecture models.

Respondents are welcome to include responses to these questions where relevant to their proposal, but they are intended as guidance and should not be interpreted as necessary.

5.A. Architecture and General System

- Does your model of Zero Trust assume service from the Cloud, or does it entail a stand-alone offering that the City would support? Is there a hybrid architecture solution that could be possible?
- What kind of use of City assets would you assume within your proposed solution? Please share your ideas for any specific network assets or other resources, under what conditions, you assume will be useful within your proposal.
- Has the system been evaluated through any third party attestations?
- How does your model make allowances for legacy applications and protocols? How would you handle the security of non-web applications?
- What is the distribution of your edge locations/points of presence (PoPs), and what kind of physical infrastructure providers do you rely on?
- How does your model address service account management and auditability of your systems?
- How do you handle encryption in your Zero Trust model, and what versions of Transport Layer Security (TLS) are allowed?
- What are the limitations of your model (total applications supported, users, etc.) in a standard licensing agreement, and what is the typical set-up of those licenses? How does it apply to external users?
- How would you both manage the resources and the financial agreement for additional capacity or heavy usage scenarios above the standard license (by user, workflow, bandwidth, etc.)?
- How do you approach application security and encryption of inbound and outbound connections?

5.B. Risk-Based Authentication/Authorization and Identity Management

- What authentication standards do you use in your model? Please help us understand what kind of policy engine/authentication process or trust broker you would use in a Zero Trust model. Do you support single packet authorization (SPA)?
- Does your policy authenticator/trust broker continue to monitor the data path after initial approval?

- How does your system integrate with other identity management providers or services? How do you handle FIDO2 compliant MFA?
- What kind of analytics do you perform on identities and users (e.g. identifying unusual activity, matching to historical patterns, etc.)?
- How do you handle multiple identity environments?

5.C. Risk-Based Endpoint Management

- Does your model focus on endpoint-initiated sessions or service-initiated sessions for applications and users?
- What operating systems, mobile devices, and endpoint agents are able to work within your model? How do you integrate or handle systems from other vendors or legacy devices? If you have integrated your model with common endpoint technology or platform providers, please elaborate.
- Are you able to assess endpoint security in a decentralized system or on unmanaged devices? Please describe how this would work within your Zero Trust model.
- Does your authentication and authorization process look at device health and security directly, or rely on any secondary products or vendors? Please describe if and how you have partnered with other providers, if relevant.
- What is your approach to routing behavior for connecting end-user devices?

5.D. Monitoring and Continuous Improvement

- Do you conduct inspections or monitor traffic streams or content for unusual or inappropriate behavior, and if so, along which parameters and in what way?
- How do you ensure an optimal user experience within an architecture that defaults to “no”?
- What is your approach in hunting product vulnerabilities? What is your policy on disclosing the results?
- How would the data on authentication flows or other types of monitoring be used to improve your risk assessments in your authentication model? Do you employ any artificial intelligence within your model?

5.E. Failure Management and Correction

- What is your approach to physical and geographic redundancy or PoPs? How does your model minimize business disruption and latency?
- Once something malicious is detected, what is your investigation and remediation approach?
- What is your model’s sequence of responses to compromised points of failure, and how flexible is that response (e.g. to the level of specific applications)?
- Who monitors for failure within your model? How is that information communicated to your clients?

6. Information and Resources

NYC3 Homepage

OneNYC - Includes City policies and previous initiatives.

NIST Special Publication 800-207: Zero Trust Architecture - Industry standard for ZTA.

Technical Vendor Resources - Includes current policies, guidelines, and standards applicable to technology projects for the City.

NYC3 and the COVID-19 Pandemic

Map of public facilities across the City - Users should filter property type by “City Owned” properties when viewing the map.