

MITRE ON THE MAINFRAME, PART 1: RECONNAISSANCE



This is Part 1 of the MITRE on the Mainframe series. For each part, we will illustrate a tactic of the MITRE ATT&CK framework in mainframe terms and technology. **The goal of this series is two-fold: 1) to emphasize that mainframe security is very relevant to contemporary enterprise security, and 2) to review mainframe security within an industry-recognized security framework.** In other words, let's make mainframe security more accessible, understandable, and relatable. To learn more about the MITRE framework itself, check out Grant McDonald's [BMC blog covering the specific tactics, techniques, and goals of MITRE ATT&CK](#).

NOTE: This blog would not be possible without the research of mainframe security experts such as Phil Young, Chad Rikansrud, the SANS Internet Storm Center, and the security vendors referenced herein. It is being shared here to encourage additional exploration of mainframe security.

During my time in incident response and security operations, I can count the number of times I heard about the mainframe on one hand.

Why mention this? While the mainframe is an important, regular topic at mainframe-focused companies, it is not as mainstream in the contemporary security space, which is dominated by distributed and cloud platforms. Think of security operations centers (SOCs) of all sizes. They are likely monitoring devices ranging from network to endpoint, data storage, etc. Yet how many have given serious thought to "big iron?"

One would think that mainframes would be part of the greater security conversation in most large organizations, at least to the extent of "should we be integrating the mainframe more?" However,

based on my observations and experience within security operations, this is often not the case. What is ironic is that the mainframe qualifies as all the device types above—a data warehouse, a high-performance computing center, and a network hub—yet it is rarely considered as equally important.

So why is this and how can we overcome it? There is ample thought leadership about this question and mainframe-specific security challenges, both operationally and perceptually. Here are a few recommendations:

[Why zOS Mainframe Security Matters](#) by Chad Rikansrud

[The Missing Link in Your SOC: Secure the Mainframe](#) by Christopher Perry

[Securing the Mainframe: How Companies can Empower Security Analysts to Protect the Backbone of Their Enterprise](#) by Christopher Perry

Without question, the lack of mainframe inclusion in security considerations is a layered issue and I encourage everyone in the security space to consider its causes. There is certainly not a simple answer, but the nuances of the challenge can help those of us in security be more prepared.

With that said, let's take a look at the relevance of mainframe security and review it within the MITRE ATT&CK framework.

Part 1 – Reconnaissance

In a previous life, if you told me: "Eddie, we need you to monitor our enterprise's most critical server. In fact, we can't even afford to have bad actors know where to look for it. Go secure it."

I would likely respond with: "Absolutely. So... *where is it?*"

Yes, it is probably in the COLO (co-location center, or [data center](#) where equipment, space, and bandwidth are available for rental to retail customers). However, this is also the first question an adversary will ask. Irrespective of whether the target is a mainframe, a public facing server, or even a network connected peripheral, any adversary must first know where to look. This is where the tactic of Reconnaissance comes into play and why it is the first tactic covered in the [MITRE ATT&CK Framework Matrix for Enterprise](#).

NOTE: MITRE ATT&CK Tactics are **not** sequential, meaning just because Reconnaissance is categorized first does not mean it necessarily happens first in practice. For example, an adversary might first gain access to credentials (*Credential Access*) then proceed to conduct Reconnaissance on the internal network. Or they may conduct internal recon again after escalating their privileges, granting them access into more sensitive areas of the network.

Per [MITRE](#):

*"Reconnaissance consists of techniques that involve adversaries **actively or passively** gathering information that can be used to support targeting. Such information may include details of the **victim organization, infrastructure, or staff/personnel**. This information can be leveraged by the adversary to aid in other phases of the adversary lifecycle, such as using gathered information to plan and execute Initial Access, to scope and prioritize post-compromise objectives, or to drive and lead further Reconnaissance efforts."*

Now let's discuss some of the Reconnaissance tactics in the context of the mainframe.

Search Open Websites/Domains


"Mainframes are searchable on the internet? Sounds unlikely."

A free and useful reconnaissance tool for searching the Internet of Things (IoT) is Shodan. Shodan is a [search engine](#) that lets the user find specific types of computers ([webcams](#), [routers](#), [servers](#), etc.) connected to the [internet](#) using a variety of filters. Some have also described it as a search engine of [service banners](#), which are [metadata](#) that the [server](#) sends back to the client. This can be information about the server software, what options the service supports, a welcome message or anything else that the client can find out before interacting with the server. Shodan collects data mostly on web servers ([HTTP/HTTPS](#) – ports 80, 8080, 443, 8443), but also collects on mainframe relevant ports such as [FTP](#) (port 21), [SSH](#) (port 22), and [Telnet](#) (port 23, 992, 1023, and 2323).


According to a [July 2017 TrendMicro report](#) on publicly exposed mainframes:

"Shodan data showed that legacy mainframes, particularly the IBM OS/390 (zSeries) and IBM AS/400 (iSeries) as well as their corresponding FTP and web servers, were exposed. The data also revealed that telecoms and ISPs were the industries with the most number of exposed/online legacy mainframes."

For illustration, the report also provided figures on types of systems exposed and geographic breakdown:



Exposed mainframe and related services by model (Credit: [TrendMicro](#))



Top countries with exposed mainframes and related services (Credit: [TrendMicro](#))

```
MSG10 OE/390
WELCOME TO THE T3 MAINFRAME AT THE CITY OF ██████████
          IBM OPERATING SYSTEM OS/390

ENTER DESIRED APPLICATION: █

YOUR IP ADDRESS: ██████████ YOUR TELNET PORT: 09100
-----LAST COMMAND:
LU: TCP10036 SENSE CODE: DATE: 05/02/17 TIME: 10:37:46
```

Exposed zSeries terminal (Credit: [TrendMicro](#))

Knowing the above, an adversary could learn [Shodan filters](#) in minutes and start reconnaissance against public facing mainframes. Following Reconnaissance, they could probe listening services to try to gain unauthorized access using common mainframe services like FTP or Telnet.

For example, unsecure FTP can serve as an entry point from which attackers can gain authorized

access to mainframes. FTP is typically used to upload transaction commands in the mainframe, such as Job Control Language (JCL), which is equivalent to a bash script in Linux® or batch file in Windows, and Restructure Extended Executor (REXX), which is similar to Python. If an FTP is successfully compromised and accessed, attackers can upload a malicious JCL program. If given executable permission to run it and with enough privileges in the terminal, it can execute arbitrary commands to [compromise a mainframe](#).

All of this just from knowing where to look!

In addition, mainframe security expert Philip Young (aka Soldier of FORTRAN) discovered that using specific keywords, not just filters, and selectively modifying them could yield further results. This is due to the specific text that HTTP banners might display being different from what one might expect. To learn more about how Shodan uses [banner grabbing](#) to find public mainframes, check out Young's blog post, [Finding Mainframes on Shodan](#).

Active Scanning (Scanning IP Blocks/Vulnerability Scanning)

An essential tool in any security professional's kit is [Nmap](#), a free and open-source utility for network discovery and security auditing. Using a simple and intuitive command line interface, users can leverage Nmap to scan specific IP addresses and CIDR blocks for open ports, listening services, and even specific vulnerabilities. While not a dedicated vulnerability assessment tool, this capability speaks to Nmap's versatility. Not only that, Nmap leverages the Lua scripting language for a variety of security use cases—including Reconnaissance of the mainframe! (For those who prefer a graphical user interface (GUI), Nmap also has a nice overlay called "[Zenmap](#)" that is quite popular.)

For example, Young has written Lua scripts that enumerate IBM® CICS® transaction IDs, CICS User IDs, NJE (Network Job Entry) target node names, TN3270 hidden fields, TSO accounts, VTAM application IDs, and even screen grab TN3270 banners, all using [Nmap](#). Users can even go beyond Reconnaissance and use Nmap for brute force attacks against NJE and TSO. These tools and more are all publicly available at Young's github repository and have been instrumental in building awareness around [mainframe security](#).

Even prominent security vendor SANS has demonstrated the availability of legacy mainframes on the internet. Without any custom scripts, the SANS Internet Storm Center leveraged Nmap to identify Telnet and SSL Telnet ports. SANS discovered that mainframe IBM® z/OS® hosts are well fingerprinted by Nmap and, though they are often labeled as OS/390, were still [very findable](#).

If you're looking for more reconnaissance tools to experiment with, another powerful tool is [masscan](#). Creator Robert David Graham claims that masscan can scan the entire internet in under five minutes, transmitting 10 million packets per second, from a single machine. It goes without saying that if you would like to experiment with tools such as Nmap or masscan, please do so in a lab environment and never scan IP blocks indiscriminately.

One important note is that Nmap is not just an external scanning tool but is equally useful for internal scanning once an adversary has a foothold in the network. Therefore, knowing network trust relationships (to be discussed below) and which systems might serve as potential pivot points for an adversary also becomes crucial in preventing effective internal reconnaissance by an adversary.

Gather Victim Host/Identity/Network/Org Information

At this point you might be thinking, "All of the above sounds interesting, but surely *our* mainframe would never be exposed publicly. We have nothing to worry about!"

Unfortunately, reconnaissance goes a bit deeper than just the "low hanging fruit" of publicly exposed systems. As MITRE aptly points out, much of reconnaissance is passive information gathering that may never involve a tool such as Nmap or Shodan. That is, it involves gathering host, identity, network, and organization information.

Here is the good news: users can make a difference in hardening the organization's security posture for each of the above categories.

Host: Practicing good cyber hygiene, such as keeping software, firmware, and configurations of access tools that may involve the mainframe (directly or indirectly) up to date, will prevent an adversary from potentially exploiting a vulnerability found in one of these resources. Application whitelisting and security configuration management (SCM) enforced at the enterprise level are very effective ways to ensure secure baselines are maintained across the enterprise. While this may not prevent zero-day attacks, it will certainly reduce host attack surface significantly. And it is often the pivot (see: Lateral Movement in MITRE) from compromised non-mainframe systems that present the greatest threat.

Identity: Securing potentially sensitive information such as personal e-mail addresses, employee PII (personally identifiable information), and credentials will also reduce any potential attack surface. Keeping this information protected as much as possible inside an internal or corporate network, ideally with multifactor authentication, will make reconnaissance more difficult for an attacker. Regarding credential security, an increasingly common practice in enterprises today is the use of password managers (or password "vaults") that store and organize usernames and passwords while offering security capabilities like the generation of complex passwords unique to each account, secure credential exchange, and real-time alerts when old or weak credentials are [exposed online](#).

Network: While more mature enterprises may not be overly concerned about public-facing mainframes, they should be concerned about the potential for "insider threats" and lateral movement (to be covered in a future post) to the mainframe once an attacker is inside the environment. Most, if not all, mainframe users today are using terminal emulation from a distributed platform such as Windows to log into the mainframe. One way to ensure that internal reconnaissance will yield limited results is to map network trust relationships in Active Directory (AD) environments using tools like [Bloodhound](#). The better one knows their internal network, the more they can prevent an attacker from using these tools to identify privilege escalation paths. With a privileged account, an attacker can likely gain additional insight into the environment or, at worst, discover a potential attack vector to the mainframe.

Organization: Adversaries will often probe open-source intelligence to find potential high value targets for social engineering or phishing attacks. While it is certainly understandable to post one's generic job title on LinkedIn, it would not be wise to go into any technical detail regarding scope of work or types of systems accessed. The unfortunate reality of the world today is that all of our public-facing activities become reconnaissance assets for adversaries. While it is not practical to stay out of public life, every user should exercise a "[zero trust](#)" posture when communicating with unknown parties and uploading data of any kind, regardless of how trusted the recipient may be.

The mainframe has long benefitted from "security by obscurity." However, as the complexity of IT

ecosystems increases, it is not as if the mainframe is simply being left out of the operational equation. If anything, it has become more crucial than ever as online transactions and the need for lightning-fast and reliable data processing continues to increase in an interconnected world. Therefore, [increased discussion of mainframe security](#) in the security community overall is needed. One way to promote this is by making mainframe security more accessible, understandable, and relatable for everyone.

This is precisely what the MITRE for Mainframe series aims to do. For our next MITRE for the Mainframe post, we'll review Initial Access and how an attacker can gain a foothold on your mainframe.