2019 Mid-market Threat and Incident Response Report

Infocyte

Q2 2019 Mid-market Threat and Incident Response Report

by Chris Gerritz, Co-founder and Chief Product Officer, Infocyte

Executive Summary

In the 90-day span from April to June 2019, Infocyte's SaaS threat detection and incident response platform, Infocyte HUNT[1], has performed over 550,000 forensic inspections on systems across hundreds of customer networks within the mid-enterprise business sector. These inspections were divided between partner-led investigations, proactive compromise assessments, and continuous use by Infocyte subscribed customers. As part of these inspections, we performed analysis on over 12.4 million unique files and 44,800 fileless in-memory injections found in whitelisted/approved applications. We reviewed 339,000+ accounts and associated behavioral logs for malicious activity, and evaluated 161,000+ unique applications for possible threats and vulnerabilities.

This report is a summary of our findings from customers and partners. It includes what we have learned about the threats and our recommendations on what companies can do to lower their risk of being affected by these and other security threats.

The findings and recommendations in this report are significant for companies in the mid-market range having between 99 and 5,000 employees and up to $1 billion in annual revenue. Most industry reports of this nature – such as the annual Verizon Data Breach Incident Report and the FireEye Mandiant M-Trends report – are focused on the large enterprise, consisting of companies that typically have far larger security budgets, larger teams and more tools and resources compared to mid-market companies. This report calls attention to the threats targeting mid-market companies and provides recommendations for those companies with more restrictive budgets and resources at their disposal.

1: Infocyte HUNT is an independent, cloud-deployable proactive security and incident response platform that helps organizations detect and respond to threats and vulnerabilities hiding within their environment. Security teams rely on Infocyte for proactive threat and vulnerability detection, on-demand incident response, and more.

Table of Contents

Key Takeaways

Ransomware – At this time, ransomware attacks are a significant threat to all types of organizations around the world. Such attacks block access to a computer system or data, usually with encryption, until the victim pays a fee to the attacker. If the ransom isn't paid, the data may be gone forever.

Recent notorious ransomware cases include the City of Atlanta, Louisville Regional Airport, Olean Medical Group, Lake City in Florida, and Onslow Water and Sewer Authority. The City of Baltimore estimates the ransomware attack they suffered will cost $18.2MM in lost or delayed revenue and restoration efforts. Managed Service Providers (MSPs) are becoming targets for ransomware as well as criminals look to infect all of the MSPs clients in a single attack.

- Of the networks assessed using Infocyte, 22% had encountered a ransomware attack that successfully executed despite endpoint, firewall and other preventative controls.

Dwell Time – The amount of time it takes for an organization to discover a threat in their environment and to remove it is known as dwell time. This metric represents a key factor influencing the overall cost and impact of a data breach when a breach occurs. For the most part, dwell time has been improving for U.S.-based large enterprises to less than 100 days. From our data, using a different methodology than most, we found dwell time of threats varies greatly by the type of threat found and is also a more significant problem for small and mid-sized organizations.

Dwell Time by Threat Type

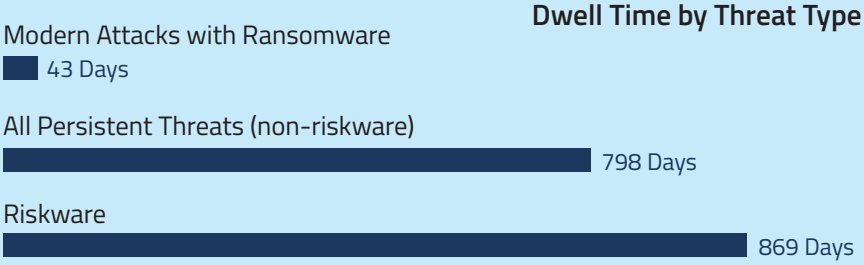

- The average dwell time for confirmed non-riskware persistent threats[2] in mid-market organizations is 798 days—far in excess of the reported times in other industry threat reports.
- Some of the longest dwelling infections (of 5 to 10 years) are of known families of malware that many EPP vendors identify as "severe" threats.
- Incidents involving modern threats like Emotet and Trickbot have much shorter dwell times, averaging 43 days from initial infection to remediation, but mostly because ransomware is also involved.

2: We define "threats" as malware flagged by at least 5 reputable intelligence sources or anti-malware engines and filtering out unwanted applications, nuisances, riskware and false positives.

Fileless malware – Memory (or code) injection is a common fileless technique used to execute external malicious code inside the private memory of another whitelisted process. Memory-only threats tend to be elusive to monitoring and protection software, as the suspicious behavior produced maps to a legitimate whitelisted process. As one of the few vendors to provide scalable volatile memory analysis capabilities, Infocyte has a unique view into the use of fileless memory-resident threats. The data provided in this report helps characterize where we're finding these threats and will benefit those hunting and responding to threats in volatile memory as well as hopefully increase the practices of this type of analysis during proactive hunts and investigations.

- Roughly 6% of the threats found were completely memory resident (fileless) -- these represented the most advanced threats discovered.

Network Hygiene – Computer systems often have unwanted software on them, some of which is preloaded when the computer is first shipped or that is downloaded and installed as part of another application's installation or use. Dubbed "riskware," this software includes dual-use administrative tools, web trackers, web toolbars and freeware versions of legitimate software bundled with adware or web activity trackers. This software is often viewed as a low priority risk to data privacy. Many companies simply ignore this software and accept the risks that come with it, rather than taking the time to eradicate the riskware to keep a clean environment. Opposing this assumption of low risk, Infocyte finds a strong correlation between the presence and dwell time of these low risk threats and the organization's overall readiness and ability to respond to a major hacking incident when one occurs.

- 63% of threats identified were classified as riskware, unwanted software, or potentially dangerous hacking or admin tools.
- 72% of our customers had multiple unwanted applications in their environment that have been there 90 days or longer (riskware dwell time)

Administrative accounts – Infocyte found a correlation between the proliferation of accounts with elevated (administrative) privileges and exposure to "living-off-the-land" lateral movement and attack patterns. We find organizations rife with issues grant far too many unnecessary privileges to employees and contractors, resulting in "dormant" but active accounts attackers may be able to compromise undetected.

- Of the systems reviewed during this reporting period, 28% of the domain accounts had elevated privileges to at least 1 system.

Methodology

During Q2, Infocyte HUNT was used to inspect 582,219 systems (mostly Linux and Microsoft Windows-based endpoints and server workloads) across hundreds of customer networks within the mid-enterprise business sector. These networks range in size between 50 and 15,000 systems, with a majority ranging between 200 and 2,000 systems. As part of these inspections, we performed analysis on over 12.4 million unique files and 44,800 fileless in-memory injections found in whitelisted/approved applications. We reviewed 339,000+ accounts and associated behavioral logs for malicious activity, and evaluated over 161,000 unique applications for possible threats and vulnerabilities.

The aggregated and anonymized data used in our analysis comes from an array of mid-market partner and customer organizations that use Infocyte's products and services. Over half of the networks that our technology is deployed into are part of a short-term engagement consisting of either:

- a proactive threat assessment (also known as a "compromise assessment") where threats may not initially be suspected or confirmed
- an incident response engagement where a security breach is confirmed and our software is deployed to assist in the investigation

The remaining systems analyzed for this report are networks that use Infocyte HUNT on a continuous basis as a proactive security capability and incident response platform as needed.

How threats were found

Within the quarter, Infocyte aggregated 32,500+ alerts of active, malicious threats and artifacts. As Infocyte is not a preventative product, these are not attack attempts but threats that made it through established defensive and preventative controls and either have artifacts of successful execution or are currently active and persisting in a network.

- 31.6% of the threats found were validated as malicious using multiple threat intelligence sources or indicators/signatures. Ransomware families such as Ryuk and Trojans like Emotet, Trickbot, Fuery, MereTam and Sality were discovered to have successfully evaded detection by the defense solution(s) customers had in place.

Expertise is critical. The rise of generalized machine learning and behavioral categorization is making it harder to characterize risk for organizations without threat analysis expertise.

- 62.8% of threats identified were classified as riskware, unwanted software, or potentially dangerous hacking or admin tools. These were flagged as web trackers, web toolbars or freeware versions of legitimate software bundled with adware or web activity trackers.
- 4.6% of threats were unknown to the threat intelligence community and had no reputation or threat intel matches but were identified by a customer or partner using our product in a proactive threat hunting capacity. A top method of discovery was finding memory-resident threats injected in core operating system processes. These flagged injections passed as clean by every analysis suite and detection engine but were discovered by Infocyte HUNT through deep host inspection, memory mapping, and analysis.
- ~1% of threats were unknown to everyone and passed through preventive security solutions as benign but were uncovered/ classified by Infocyte analysts supporting our customers and partners.

Types of Threats Found

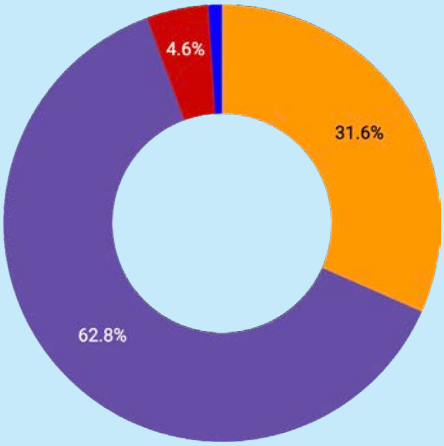

31.6%	Known malicious that had bypassed defense solutions
62.8%	Riskware, unwanted software, or potentially dangerous hacking or admin tools
4.6%	Unknown to the intelligence community
~1%	Unknown to everyone

61% of all detections of active (non-riskware) malware are made with a generic detection such as a heuristic or behavior signature or machine learning categorization algorithm. These detections often require additional verification to confirm and make it difficult to measure and communicate risks to the business. The rest had specific signatures naming the threat or malware family.

"Despite talk about signatures being 'dead' [as a detection method], many security analysts continue to rely on them — when they do work, they name the threat," says Chris Gerritz, Co-founder and Chief Product Officer, Infocyte.

"We find more specific ML categorization enriched with threat intelligence that can name and classify the threat or risk makes advanced detection much easier to consume."

Findings

Among the findings from the system inspections are the following:

- 22% of scanned networks have encountered ransomware attacks in the last 90 days that went undetected by their anti-virus, endpoint security tools, and other preventative controls.
- 5% of networks have experienced an advanced targeted or multi-stage attack this year which involved an initial adversary entry vector followed by deployment of additional malware and hacking tools like Mimikatz plus use of stolen credentials.
- 72% of environments had multiple unwanted applications (dubbed "riskware") in their environment that have been there 90 days or longer (riskware dwell time) which we correlate as an indicator of general hygiene and control of the network.

Riskware correlates to poor network hygiene. Despite the low priority of threats that "riskware" encompasses, Infocyte finds that those organizations that struggle with controlling risky and/or unwanted applications have a lower readiness to respond to priority threats when they are found.

Mid-sized enterprises struggle with dwell time

Multiple industry reports including the Verizon DBIR, FireEye Mandiant M-Trends and others report on the time it takes organizations to discover threats in their environment and remediate them. This metric is called dwell time and it represents a key factor in the overall cost of a data breach when a breach occurs. For the most part, this metric has been improving for U.S.-based large enterprises. Nevertheless, our data shows that dwell time varies widely depending on the type of threat and is still a large problem for small and mid-sized organizations.

Typical dwell time numbers reported by a majority of industry reports are generally representative of confirmed reportable breaches (confirmed information disclosure or damage). Using the forensic data Infocyte collects from systems, we are able to measure dwell time[3] for all active persistent threats and malware found during inspection that may or may not have a reportable business impact but that potentially compromises the integrity, confidentiality or availability of an information asset. This viewpoint paints a very different picture from most industry reports:

- The average internal dwell time for confirmed persistent threats (i.e. malicious malware that is confirmed by at least 5 reputable sources, and after filtering out hack tools, riskware, adware, and unauthorized applications as well as a concerted effort to remove false positives) is 798 days—far in excess of the reported dwell times in other industry reporting.

Mid-sized companies struggle with dwell time. Persistent malicious threats can dwell on networks for years – sometimes as many as 5 to 10 years – without detection by traditional means.

3: Infocyte calculates dwell time as a combination of time-to-detect a threat and time to respond/remove the threat from the network. These are calculated from measurements such as the earliest timestamp (i.e. file creation time on a disk) of an artifact, the time Infocyte first detected the artifact, and the timestamp of the latest occurrence ("Last Seen").

- 54% of the networks scanned were discovered to have threats (non-riskware) dwelling longer than 30 days.
- While the type or severity of threat did influence dwell time findings, how the discovery was made, whether during a compromise assessment or an incident response engagement, had no observable impact. In fact, some of the timestamps involved in these incidents were so long ago (e.g., 9 years ago) that we initially assumed it must be tainted data or that timestomping was occurring. While timestomping did, indeed, occur in a couple of these incidents, there is a significant amount of correlation with multiple timestamp sources that indicates that these threats really did survive for years.
- Some of the longest dwelling infections (5-10 years) were of Neshta, Fluery, and other relatively known malware families originating from a previous decade but still listed by EPP vendors that detect them as "SEVERE." We have evidence of these being active since 2009 on some older systems.
- Among just the analyst-confirmed Trickbot infections (n=75), the average dwell time was 43 days. Like Emotet, Trickbot is a modern family of remote access tool (RAT) type malware that has active groups behind it, and first appearing around 2016. It has been a highly ranked threat as of 2018.

Most Trickbot attacks had significantly lower dwell times because malware like Trickbot was often used to deploy ransomware (with gaps of days/weeks between) which alerted the victim of the compromise. This, in turn, began the response process which included engaging Infocyte-armed responders.

This gap between initial trojan (trickbot) infection and the ransom is a hallmark of modern targeted ransomware cases over the last 12 months. Trickbot was usually detected by Infocyte during the ransomware response process and was rarely the initial indicator of a security incident.

There were also several examples where Trickbot was detected and mitigated before damage occurred (0 days of dwell time).

A possible explanation for some of the extreme dwell times we measured could be that many of the active infections residing on these systems are beaconing to sinkholed domains and posed no immediate threat. We found this in several cases we reviewed. One family of Infections we traced back to 2009-2011 on multiple networks posed no continuing threat after a series of botnet operators were arrested in subsequent years. Still, it was surprising to find the malware still active on what appear to be protected endpoints so many years later.

Unfortunately, due to the timespan, we have little hope in determining what the criminals were able to do or steal prior to their command and control networks being brought down.

Another set of examples exhibiting long dwell time appeared to be systems that were part of untargeted botnets. We assume the attackers simply utilized these systems for targeting other organizations and never posed a major threat to the infected organization itself. Ultimately, the organizations had no perceived business impact from the infections which further contributed to long dwell time.

Dwell time of low priority threats is alarming

We found the dwell time for lower priority threats (i.e., riskware, adware, unwanted applications, etc.) to be much longer than severe threats, averaging 869 days of dwell time (with statistically relevant upper ranges approaching 7 years). These numbers indicate that these threats are either being ignored or simply happen with such frequency that organizations can't keep up with the flood of unwanted, unauthorized, and risky applications.

- On average, 11 distinct types of riskware was discovered per network scanned, ranging from adware to bitcoin miners installed by employees.
- The percentage of networks with unwanted applications lasting longer than 90 days is 72%. Among these, the average number of distinct threats (different file hash) is 6.
- While riskware, unwanted applications, sinkholed malware, and otherwise untargeted botnet membership pose relatively low risk to the infected organization and are often overlooked, Infocyte finds a strong correlation between the presence and dwell time of these threats and the organization's overall readiness and ability to respond to a major hacking incident when one occurs. In other words, a clean environment is an indicator of positive control and control of a network is necessary when responding to an active, spreading infection.

A clean environment is an indicator of positive control and control of a network is necessary when responding to an active, spreading infection.

Fileless malware continues to grow

Memory (or code) Injection is a fileless technique used to execute external malicious code inside the private memory of another whitelisted process. Detecting malicious fileless injection events behaviorally, as some of the more advanced endpoint protection products do, generally relies on hooking critical operating system functions which can slow systems down and/or cause a lot of false positives. Infocyte's methodology in analyzing these threats is different from protection tools as we perform live forensic analysis of memory to identify and extract rogue executable code not mapped to a file on disk. This gives us a unique view of these types of attacks and copies of the resulting executable payloads they place in memory, granting much higher fidelity and signal-to-noise.

As one of the few vendors to do this type of analysis at scale, we have a unique view into the use of fileless memory-resident threats as well as the examples of legitimate use of this operating system feature. Our hope is that this data will benefit those hunting and responding to threats in volatile memory as well as increase the practice of this type of analysis during proactive hunts and investigations.

Here are a few observations and figures regarding the use of these techniques:

- Memory-resident fileless attacks leveraging code injection are common to nearly most modern threat actors and penetration testers—it's not just for advanced persistent threats (APTs) anymore.
- Most modern malware families like Trickbot, Emotet and Ryuk all leverage fileless injection techniques but they are not completely fileless. Memory-resident portions are either one component of the malware, a stage in the attack, or the final payload left after cleaning up other stages.
- Adversaries leveraging stolen admin credentials and powershell-based "living-off-the-land" techniques will often still utilize in-memory code injection to load various components they need to conduct their actions.
- Memory-only threats tend to be elusive to monitoring and protection software as the suspicious behavior produced maps to a legitimate, whitelisted process.

- Code injection is not considered a vulnerability or flaw of the operating system, but in fact, a feature used by legitimate software. It works in all modern operating systems including OSX and Linux (as does Infocyte). Linux and OSX injection is significantly more rare as these platforms don't have protection software requiring that level of evasion.
- As traditional file-based analysis can sometimes fail against injects, the most important factor in determining whether injected code is malicious or benign is the location or process it is injected into.
  - Most legitimate memory injection is inside .NET and Java applications performing their own in-memory Just-in-Time (JIT) compilation. Software from Oracle, EMC, etc. written in these languages are the most common sources.
  - Core Operating System executables are almost never the targets of legitimate code injection unless it's another security product. These should always be viewed with suspicion.
  - 61% of injections we found in critical Microsoft Windows processes were found to be malicious, the rest were security tools. The benign injections found (39%) were almost exclusively done by legacy security products monitoring other processes and are fairly easy to filter out. This type of injection is not performed by modern, kernel-mode protection software.
- Memory protections assigned to malicious code are another important factor in determining whether code is malicious or benign. 76% of malicious injections use the most liberal Read/Write/Execute memory protections (i.e., no protections). The rest apply Read/Execute permissions following the injection event which helps hide it more among legitimate code injections.

Top 5 memory inject locations for confirmed attacks (n=2,056)

1. chrome.exe - 31%
2. iexplore.exe - 15%
3. svchost.exe - 8%
4. powershell.exe - 8%
5. regsvcs.exe - 8%

- Average size of memory injects in core Windows processes: 129kb
- Average size of malicious injections: 518kb
- Largest malicious injection: 2.9MB
- Smallest malicious injections: 4kb (a single function loader stub will easily fit in a single page of memory but never contains the full malware payload)
- Most fileless malware has multiple components of various sizes injected into the same process

Top 10 processes with any memory injection (includes benign injections) (n=45,843)

1. ctxgfx.exe - 17.29%
2. ekrn.exe - 14.50%
3. chrome.exe - 8.52%
4. psanhost.exe - 6.19%
5. driversupportao.exe - 5.95%
6. emc.captiva.webtoolkithost.exe - 4.35%
7. firefox.exe - 3.62%
8. toad.exe - 2.83%
9. svchost.exe - 2.33%
10. driversupportaosvc.exe - 2.16%

- Average size of memory injects: 212kb
- Largest size of memory inject: 22.2MB

Administrative Account usage and distribution

While evaluating exposure to potential living-off-the-land lateral movement and attack patterns, we noticed a correlation between the proliferation of accounts with elevated (administrative) privileges. In general, we find organizations with a high amount of issues give far too many unnecessary privileges to both employees and contractors.

The average Domain Administrator or equivalent privileged account (defined here as an account with administrative rights to remotely administer multiple systems and execute at SYSTEM, root, or sudo levels) had login activity on 14 systems on average. One way we detect the malicious use of privileged accounts is to baseline this number for each account and track spikes.

- Of the systems that we reviewed during this reporting period, the ratio of domain accounts with elevated privileges to unprivileged accounts averaged 28% or 1:4 (n=48,993).
- For customers with traditionally managed networks with established least privilege policies, the average ratio is closer to 10% or 1:10.

What companies can do to lower their risk

Many small to mid-sized companies may falsely assume that they are unlikely to be targets of cyber threats, and that their large enterprise counterparts are at higher risk for attack. Infocyte's analysis in this report shows just the opposite—that mid-market companies have threats lurking in their systems for significantly long amounts of time. What's more, these threats have made it past the defensive systems such as network monitoring detection and endpoint protections and now have free reign of internal systems until they are discovered and neutralized. So, what can companies do to lower their risk?

- Become aware of unknown risks. Compromise assessments (CAs) by an independent 3rd party are becoming a best practice much like the venerable penetration test has been. A competent CA vendor can independently hunt for and detect advanced threats (e.g., those that are memory-only) without having to modify your current security stack. CAs can find what your security system is missing or validate your investments and ultimately exist to help organizations understand the threats and risks to the business.

- When it comes to detecting advanced threats, don't ignore the endpoint or privileged identities. Proper endpoint and privilege usage visibility are the mainstays of modern detection and response. We find networks that rely solely on traditional firewall or network traffic monitoring are at the highest risk from long dwelling persistent threats as they are able to hide on endpoints and/or steal credentials to blend in.
- Pay attention to network hygiene and institute processes to reduce or eliminate riskware and unauthorized applications from the environment. For example, downloads of certain web tools can be blocked, and "bloatware" (that software that comes pre-installed on computers) can be removed before the computers are deployed. System hygiene is important, not only for lowering risk but also for making it easier to respond to serious security incidents. Companies that are able to maintain good network hygiene have a higher level of readiness to respond to future business impacting issues.
- The cleanest and most incident-ready networks we've worked with have more than just written policies or the visibility and tooling provided by products like Infocyte. The most effective organizations we work with also have the empowerment to act on the information: to control what software is authorized and have the backing of the business or executive leadership to enforce policy. Proper tooling to enforce policy and act on issues found is a necessary force multiplier to make small teams more effective.
- Gather metrics, baseline data and industry benchmarks to understand the organization's current security posture. Not only is this information necessary for measuring change (improvement!) over time, but senior executives will want to know this information before considering budget approvals for new or additional security solutions or staffing.

This may seem overwhelming, but there are several ways to solve these problems and be ready for an incident should one occur. Response plans are a good place to start but we find they are often unexercised in many smaller organizations. Having a good partner that can help guide you through an incident and proper tooling and visibility can make a world of difference. Organizations that lack the ability to perform modern detection and incident response functions are encouraged to find a managed detection and response (MDR) provider which specializes in this. Lastly, good network hygiene will help in identifying anomalies and although it starts with visibility tooling, it really ends with action: A written policy is not enough, there has to be consistent enforcement.

Conclusion

This data, gathered and analyzed through Q2 of 2019, is representative of a cross section of mid-market enterprises and small businesses with centrally administered networks. Using actual live views into hundreds of networks, it demonstrates the reality of threat dwell time to be much longer than that represented in other industry reports—especially for the types of threats where active business impact is not readily apparent. Additionally, there is significant value in the data we have on fileless malware and attacks for threat hunters and incident responders to tune their methodologies for finding these elusive threats.

There is still much work to do to reduce the risk and improve the readiness of small and medium sized businesses to respond to modern cyber threats. Some of this work needs to start with general network hygiene practices and establishing control of what is allowed on the corporate network.

To learn more about Infocyte, request a copy of our Mid-Market Threat and Incident Response report, or to schedule a live demo of our platform, please visit www.infocyte.com.

About the report author

Chris Gerritz, co-founder and chief product officer of Infocyte, has over a decade of experience hunting adversaries within some of the largest and most targeted defense networks in the world. From building the U.S. Military's first enterprise-scoped threat hunt team to supporting top threat assessment and incident response teams, Chris has led, participated in, or supported over a thousand of cybersecurity engagements in networks throughout the world.

About Infocyte

The world's leading cybersecurity and incident response firms (Check Point, PwC, Grant Thornton, and more) use Infocyte's platform to proactively detect and respond to vulnerabilities and threats hiding within their customers' endpoints, data centers, and cloud environments.

Large enterprises with a security operations center (SOC) leverage our platform to maintain compliance, reduce risk, and optimize security operations.

Small and mid-market organizations with an understaffed security team and fewer technical resources, leverage Infocyte as a managed service, delivered through one of our partners, providing enterprise-level detection and response services to the mid-market.

For partners, Infocyte represents the fastest path for delivering cost-effective and flexible consulting services (i.e. compromise assessments and incident response) and ongoing Managed Detection and Response (MDR) services to their customers via our easy-to-use cloud platform, Infocyte HUNT.

Infocyte

3801 N. Capital of TX Hwy
Suite D-120
Austin, TX 78746

(844) 463-6298
sales@infocyte.com