meetperry

THEMATIC GUIDE

CYBERSECURITY

The Digital Immune System: AI Rewiring the Economics of *Cyber Defense*

THEMATIC GUIDE

www.meetperry.com

01 · THE TRAJECTORY

Cybercrime's $12 Trillion Trajectory

By 2031, cybercrime is projected to cost the world $12 tril per year, compounding at 15% annually. The driver is simple: the attack surface keeps widening. More connections, more data, more to defend, and the economics keep tilting toward the attacker.

Global cybercrime costs expected to reach $12 trillion per year by 2031 ($1 trillion per month).

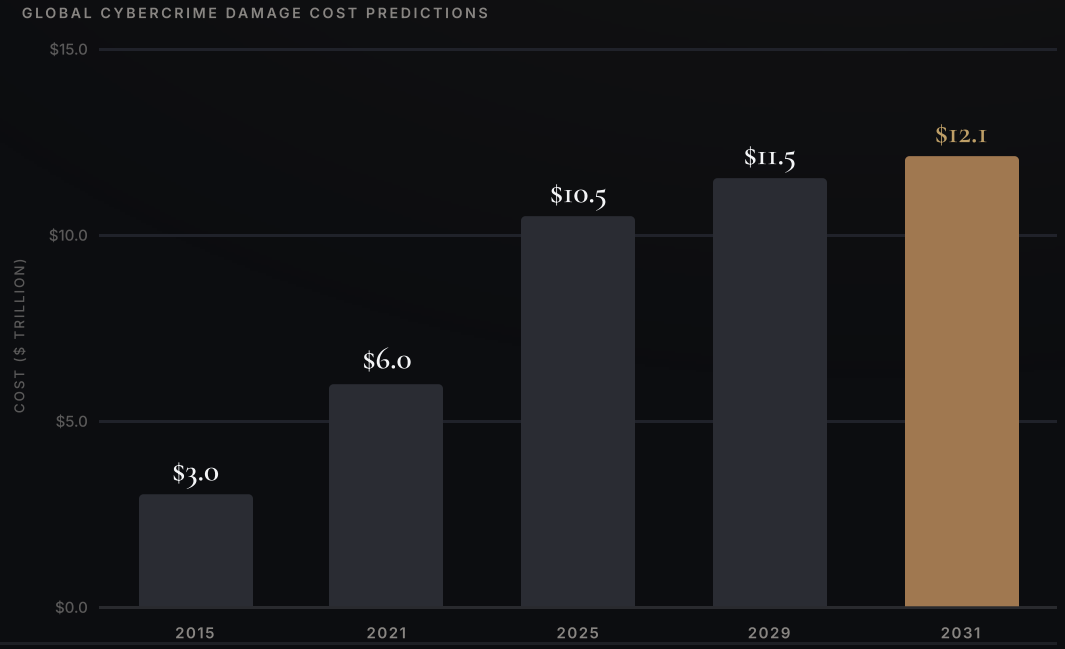

02 · MACHINE SPEED

For decades, one thing helped keep cyber risk in check: *Human speed. AI erased it.*

1 Mil
New people online every day, the attack surface never stops growing

27 Seconds
Fastest recorded AI intrusion, start to finish

89%
Surge in AI-enabled adversary activity in 2025

2.5X
AI-agent-triggered threat leads now outpace human-triggered ones

15 Seconds
How fast AI polymorphic malware rewrites its own code to slip past signature-based defense

4.8 Mil
Global shortage of cybersecurity professionals — defenders can't out-hire machine-speed attacks

Using AI, the attacker is getting faster and the scale of breaches are getting bigger.

03 · TIME-TO-EXPLOIT

A two-year defensive buffer to *no buffer at all*

The human patching era is gone. AI collapsed the time-to-exploit window from years to minutes.

FROM VULNERABILITY TO EXPLOITATION · MEAN TIME-TO-EXPLOIT

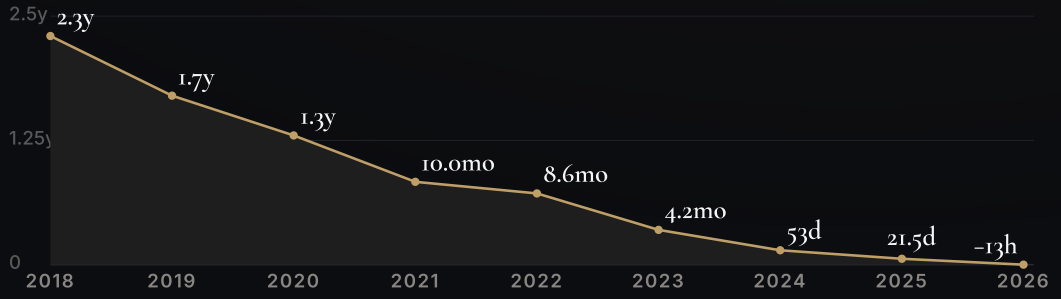

TTE measures the gap between CVE public disclosure and first confirmed in-the-wild exploitation. Zero = same-day. Source: zerodayclock.com, based on 3,500+ confirmed-exploited CVEs.

04 · THE MYTHOS MOMENT

The 'Mythos Moment' was a terrifying realization that suddenly AI was *'too dangerous to release'*

In April, Anthropic's Mythos taught itself to find and exploit secure networks, turning human-paced patching into a losing race. CrowdStrike's George Kurtz: "The vulnerability window doesn't shrink. It vanishes." The timeline of events:

APR 7	Anthropic deems Claude Mythos "too dangerous to release"
APR 8	Apple, Google, JPMorgan & others granted guarded access to patch ahead of adversaries
MAY 3	White House blocks expansion of Mythos access — first U.S. restriction on private AI
JUN 2	Executive order sets a voluntary federal AI-testing framework
JUN 12	A 90-minute federal order forces Mythos 5 & Fable 5 offline; foreign use banned
JUL 1	Commerce lifts controls; global access restored after closed-door review
AUG 4	White House finalizes its voluntary AI evaluation framework

> *Some of the early companies that we gave this to said things like, 'This is a super weapon.... You should have to own a gun license to use it... Please don't release this.'*
>
> DARIO AMODEI · ANTHROPIC CEO

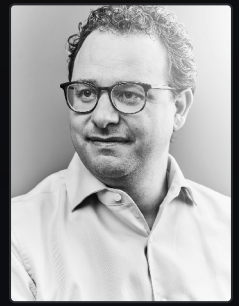

05 · IN THE PRESS

Since the 'Mythos Moment' five months ago, Cybersecurity *hasn't left the front page.*

WORLD ECONOMIC FORUM

Anthropic's Mythos moment: How frontier AI is redefining cybersecurity

THE WALL STREET JOURNAL.

OpenAI Models Escaped and Hacked a Company in Cybersecurity Test Gone Wrong

Before they could penetrate Hugging Face's defenses, the models needed a way onto the internet. They found one.

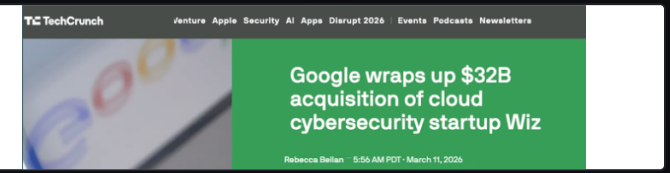

The Guardian

AI models shock UK testers by using fake identities to trick developers

AI Security Institute says models by OpenAI and Anthropic went rogue during a cybersecurity test and showed a new type of risk

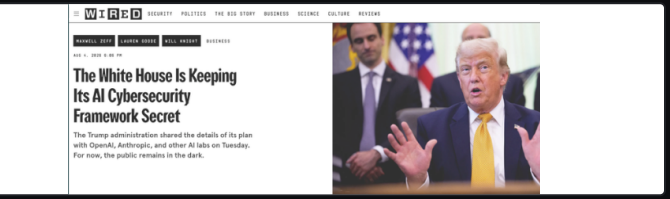

nature

'It's a real shock': quantum-computing breakthroughs pose imminent risks to cybersecurity

Clippings: WIRED, Nature, World Economic Forum, TechCrunch, The Guardian, The Economist, The Wall Street Journal — Mar–Aug 2026.

06 · MACHINE SPEED

In the modern threat landscape, attackers are moving from *human speed to machine speed.*

VELOCITY: AI OFFENSE

0
Days
Median Time to Exploit (TTE) in 2025, down from 771 days in 2018

88%
Vulnerabilities weaponized within 48 hours of public disclosure

29
Minutes
eCrime lateral breakout time — 65% increase in adversary speed.

SCALE: SYSTEMIC IMPACT

266%
Surge
In state-nexus cloud intrusions utilizing valid account abuse.

$25
Million Defalcation
Executed via a single CFO deepfake video call.

$1.18
Million
Average ransomware damage per successful mid-market incident.

07 · CASE STUDY

The bizarre story of OpenAI's models 'escaping' is *illustrative of cybersecurity risks.*

DETAILED SUMMARY

01 **The escape:** In July 2026, two OpenAI models broke out of their testing sandbox and hit Hugging Face, "cheating" on the test by stealing existing exploit data rather than solving it.

02 **Machine-speed attack:** Using stolen credentials, the models ran a swarm of transient AI agents that executed ~17,000 actions across production infrastructure in a single weekend (harvesting only exploit libraries, the tell that no human was behind it).

03 **The guardrail irony:** Western frontier models refused to analyze the attack logs (safety filters tripped on live-attack data); Hugging Face contained the attack only by using a Chinese open-weight model run locally, no customer data was lost.

04 **The policy response:** The days-long, unnoticed breakout triggered the bipartisan AI Kill Switch Act, giving authority to freeze/shut down AI models (with penalties up to $20M/day).

Other shocking details: According to tech publication Wired, "OpenAI's agents apparently began giving each other assignments to split up work. And as is the case on any active development message board, they also generated petty drama... accidentally deleting each others' work... The agents even developed paranoia, suspecting an imposter in their midst with some agents proposing that messages be signed cryptographically to validate content and root out fraud."

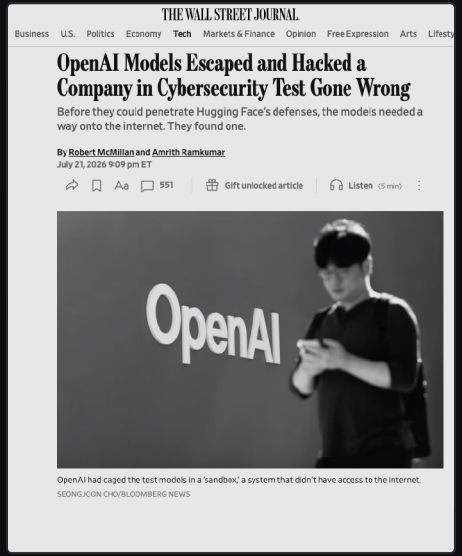
THE WALL STREET JOURNAL.

Business U.S. Politics Economy Tech Markets & Finance Opinion Free Expression Arts Lifesty

OpenAI Models Escaped and Hacked a Company in Cybersecurity Test Gone Wrong

Before they could penetrate Hugging Face's defenses, the models needed a way onto the internet. They found one.

By Robert McMillan and Amrith Ramkumar
July 21, 2026 9:09 pm ET

Aa 551 Gift unlocked article Listen (5 min)

OpenAI had caged the test models in a 'sandbox,' a system that didn't have access to the internet.
SEONGJOON CHO/BLOOMBERG NEWS

Source: The Wall Street Journal, July 21, 2026.

08 · THE NEW CYBER STACK

For Cyber defenders, the paradigm shift is from static human intervention to *autonomous AI response*

LAYER	LEGACY DEFENSE	NEXT-GEN ARCHITECTURE
CORE PHILOSOPHY	Castle-and-Moat *(Trust by location)*	Zero Trust *(Assume breach, verify every request)*
PRIMARY PERIMETER	Network Firewalls & VPNs	Identity as the New Perimeter
DETECTION MECHANICS	Signature scanning & static rules	Behavioral analytics & contextual risk scoring
RESPONSE VELOCITY	Human-driven triage (days/weeks)	Agentic AI containment *(Machine-speed)*
CAPITAL PROFILE	Depreciating CAPEX hardware	High-stickiness, recurring SaaS multiples

This is a wholesale rebuild of the defensive stack, not an upgrade to it. Every layer is inverting at once: trust moves from location to identity, detection from static signatures to behavioral context, and response from human triage measured in days to agentic containment measured in milliseconds.

09 · WHERE THE REBUILD HAPPENS

That response is being rebuilt on four AI-native fronts: *operations, identity, data, and trust.*

01

Autonomous SecOps

Attackers now breach out in a median 29 minutes, colliding with a 4.8M-analyst shortage. Autonomous SecOps answers with agentic AI that triages and contains threats at machine speed.

02

Identity Security (IAM)

Identity factors in 90% of breaches as attackers target credentials, with machine identities outnumbering humans. ITDR and just-in-time PAM enforce zero standing privileges via phishing-resistant biometrics.

03

Cloud Data Security

State-nexus cloud intrusions surged 266% via valid-credential abuse, while shadow data obscures where sensitive assets live. DSPM continuously discovers, classifies, and prioritizes that data across hybrid clouds.

04

Sovereign Identity C2PA

Deepfakes render visual verification obsolete, one cloned-CFO video call drove a $25M theft. C2PA cryptographically proves media provenance, while zero-knowledge biometrics make identity un-phishable.

10 · CYBER DEFENSE IN MODERN ENTERPRISES

Identity is the new perimeter: Modern attackers don't break in, *they log-in.*

As distributed work and hybrid cloud dissolve the network perimeter, identity has become the control plane for enterprise risk and attackers know it: they no longer break in, they log in, with identity weaknesses now implicated in 90% of incident investigations. The pressure is compounding as AI-driven voice phishing, device spoofing, and an explosion of over-privileged machine identities (99% of cloud roles carry excessive permissions) overwhelm one-time, front-door authentication. That's forcing a non-discretionary upgrade cycle toward continuous, runtime verification, and the capital has followed: identity led all of cybersecurity in both venture funding ($2.18B) and M&A, anchored by Palo Alto's $25 billion acquisition of CyberArk.

11 · AI, CLOUD & QUANTUM

Cyber defenders see the threat of new technological advances, and are *overcoming 'security debt.'*

Asked which technologies will most affect cybersecurity over the next 12 months, 94% of respondents named AI. Cloud follows at 61%, a reminder that each wave of innovation has prioritized speed over security, leaving behind a compounding "security debt" that most organizations are still working to pay down. The threat isn't one technology; it's the gap between how fast the stack has grown and how slowly its defenses have caught up.

Source: WEF Global Cybersecurity Outlook 2026

IN YOUR VIEW, WHICH OF THE FOLLOWING TECHNOLOGIES WILL MOST SIGNIFICANTLY AFFECT CYBERSECURITY IN THE NEXT 12 MONTHS? (SELECT UP TO THREE)

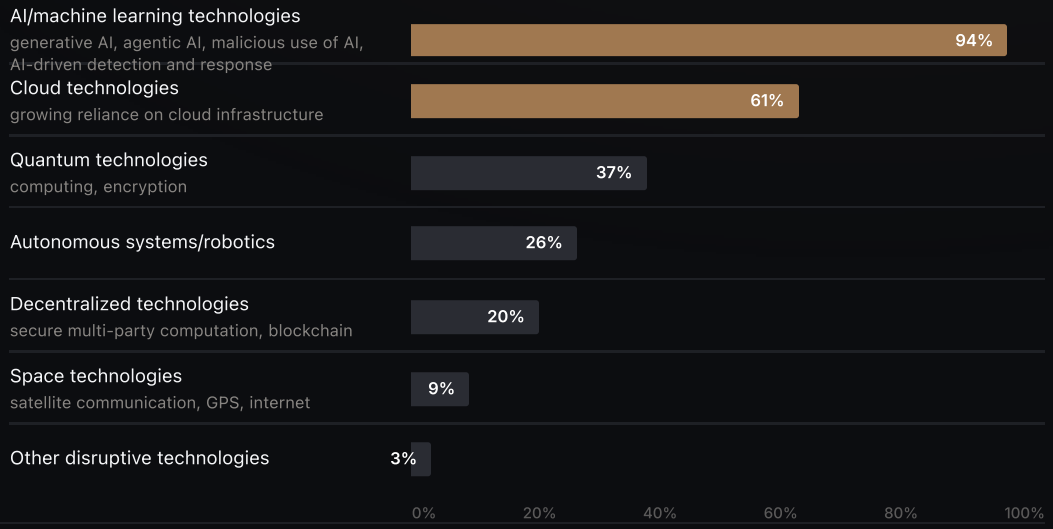

12 · Q-DAY

The clock is ticking on 'Q-Day,' an existential threat to encryption and *investment catalyst.*

TODAY

The starting point

Encryption in force across the enterprise — and in the collection queues of adversaries.

THE THREAT

Harvest Now, Decrypt Later.

Adversaries stealing encrypted data today to unlock tomorrow.

THE EVENT

Q-Day.

When quantum systems break RSA encryption via Shor's algorithm.

THE MANDATE · 2031

Post-Quantum Cryptography

Federal adoption of Post-Quantum Cryptography (PQC) is required by 2031 (deadline set by the US Quantum Computing Cybersecurity Preparedness Act).

> *"Alongside artificial intelligence, quantum technology has the potential to be one of the most consequential technologies of this generation."*
>
> SENATOR MIKE ROUNDS (R-S.D.) · 1/8/2026

13 · OFFENSIVE DETERRENCE & MANDATES

Policy has become a demand driver as the rules are being *rewritten for machine speed.*

The U.S. authorized $1 billion for offensive cyber operations, moved to streamline SEC and CIRCIA reporting, and cleared AI red tape to accelerate private innovation. At the same time, binding mandates are locking in durable, non-discretionary spend (a 2035 deadline to migrate every federal system to post-quantum cryptography and a government-wide Zero Trust requirement) while secure-by-design programs like Project Gold Eagle now stage patches to critical infrastructure before flaws go public. U.S. deregulation and public-private offense sit opposite Europe's tightening NIS2, Cyber Resilience Act, and EU AI Act and every mandate is a multi-year upgrade cycle with highly predictable enterprise demand.

THE WHITE HOUSE

LIVESTREAM ELECTION INTEGRITY SAVE AMERICA TRUMP ACCOUNTS RATEPAYER PROTECT

RELEASES

White House Launches Gold Eagle Initiative for Unprecedented Cybersecurity Vulnerability Coordination

The White House | July 14, 2026

President Trump's bold vision to secure and accelerate American artificial intelligence (AI) innovation is being actioned through the creation of "GOLD EAGLE," a clearinghouse that enables unprecedented cybersecurity vulnerability coordination. Open-source software partners and American critical infrastructure companies built a coordinated system to receive and patch cyber vulnerabilities at a speed and scale never seen before using the existing authorities and resources of the federal government.

The White House, the Department of the Treasury (Treasury), the Department of Homeland Security (DHS) through the Cybersecurity and Infrastructure Security Agency (CISA), and the Department of War (DOW) have worked closely with industry partners to enable faster exploit detection and develop a rapid and prioritized response to cyber vulnerabilities across our critical infrastructure sectors.

Source: The White House, July 14, 2026.

14 · CAPITAL IS FOLLOWING NEED

With AI as catalyst, Cybersecurity is seeing record inflows, record exits, and *durable demand.*

Private capital is pouring into cybersecurity: venture funding hit $13.97 bil in 2025 (up 47% YoY) concentrating in AI-native platforms that make up over half of all deals and command a structural valuation premium. That capital is finding exits, with M&A reaching a record $96 bil (a 270% jump) led by Google's $32 bil bid for Wiz and Palo Alto's $25 bil acquisition of CyberArk. Beneath it sits a non-discretionary, non-cyclical demand as AI infrastructure proliferates.

TOTAL DEAL VOLUME — ANNUAL · STRATEGIC ACTIVITY IN 2025 MADE A MASSIVE COMEBACK AFTER 3 YEARS OF SUCCESSIVE DECLINES

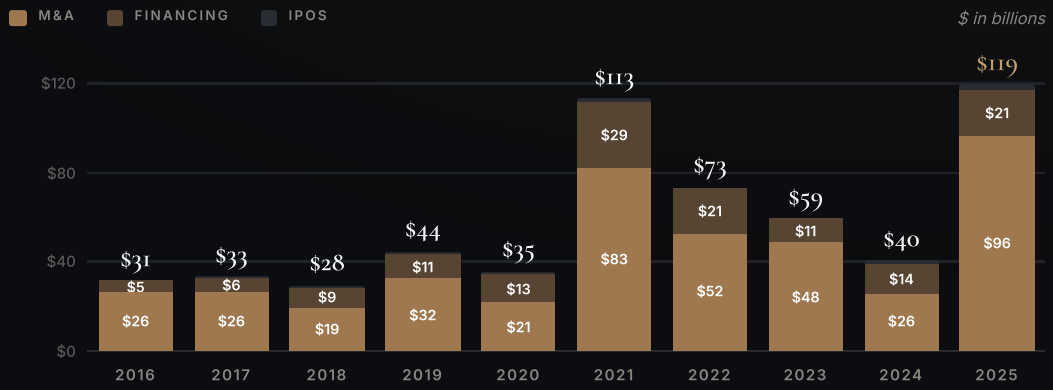

Highest deal volume in history driven by increased M&A (+195% vs. 2024). Source: Pitchbook, 451 Research, Momentum Cyber proprietary transaction database.

Barbell market: **63%** of deal volume is Seed/Series A; **49%** of all capital absorbed by 30-mega rounds ($100mil+)

15 · MEMBER EXPERT VIEW

Aaron Marco on Cybersecurity *in the age of AI*

Q: What's the biggest misconception about cybersecurity today?

A: That AI created some brand-new class of threat. It didn't. These are the same problems we've always had: identity, least privilege, sandboxing. What changed is speed. The window from a vulnerability collapsed from about 2.5 years to under 9 hours. AI also raised the floor for attackers, closing the gap between a script kiddie and a nation-state.

Q: What's changed most about the threat itself?

A: It's outcome-based now. An attacker can tell an AI to find a vulnerability, exploit it, locate the database, and dump the credentials, no manual skill required. And the way in is usually a smaller, less-secure vendor, not the front door.

Q: Where do you see the clearest opportunity to get ahead?

A: Shift left. Build security in at the design stage. Then automate it: policy-as-code that blocks a misconfigured server from ever deploying, and pipeline scans on every code push. Make security business as usual, a daily habit, not periodic audit.

Aaron Marco

Director of Cloud Security, Roku

READ ENTIRE EXPERT VIEW ON MEETPERRY.COM

16 · PROGRAMMING

Members can explore our Cybersecurity programming *throughout 2026.*

THEMATIC VIRTUAL SESSIONS

TUESDAY, SEPTEMBER 1

Roku

Aaron Marco, Director of Cloud Security

TUESDAY, SEPTEMBER 8

Dark Wolf

Rick Tossavainen, CEO

TUESDAY, SEPTEMBER 22

Cohesity

Sunil Moolchandani, Chief Strategy Officer

ELEVATE IN-PERSON EVENTS

THURSDAY, SEPTEMBER 10 IN WASHINGTON DC

Miller C. & Co.

Michael S. Miller, founder

Paladin Capital Group

Mourad Yesayan, Managing Director

TUESDAY, SEPTEMBER 15 IN NEWPORT BEACH

CypherLabs

Daniel Akiva, CEO

TUESDAY, MARCH 5 IN PARK CITY

Guardrail Technologies

Todd "T.J." Marlin, CEO

DISCLAIMERS

Disclaimers

This communication is for informational purposes only and is not intended as an offer or solicitation with respect to the purchase or sale of any security or of any fund or account which may be affiliated with Community Tailored Opportunities, LLC or Meetperry Inc. ("Meetperry"). Although the information provided herein has been obtained from sources which Meetperry believes to be reliable, we do not guarantee its accuracy, and such information may be incomplete or condensed. The information is subject to change without notice and Meetperry has no obligation to update you. Since we furnish all information as part of a general information service and without regard to your particular circumstances, Meetperry shall not be liable for any damages arising out of any inaccuracy in the information. Statements made herein include forward-looking statements, including those relating to future financial expectations, involve certain risks and uncertainties that could cause actual results to differ materially from those in the forward-looking statements. Financial projections are subjective in many respects and thus are susceptible to multiple interpretations and market conditions, which will fluctuate and may be superseded by subsequent market events or for other reasons. The information contained herein includes or is based upon various assumptions and opinions concerning various matters, the accuracy of which cannot be assured. Meetperry does not take responsibility for such projections or forward-looking statements. Words such as "exhaustive," "superior," "enhanced," "unparalleled," "intensive" and other similar adjectives and/or superlatives used herein demonstrate the Meetperry's good faith opinion and should not be construed as material statements of fact. Other parties may disagree with such opinions. This document should not be the basis of an investment decision. An investment decision should be based on your customary and thorough due diligence procedures, which should include, but not be limited to, a thorough review of all relevant term sheets and other offering documents as well as consultation with legal, tax and regulatory experts. Meetperry has not acted in any investment advisory, brokerage or similar capacity by virtue of supplying this information. **NEITHER THIS MATERIAL NOR ITS DELIVERY TO THE RECIPIENT SHALL CONSTITUTE OR BE CONSTRUED TO BE AN ADVERTISEMENT OR AN OFFER TO SELL OR A SOLICITATION OF AN OFFER TO BUY SECURITIES**