Some insights into Hardware Security Modules (HSMs) in the Cloud

davinune@microsoft.com

Microsoft’s HSM products

- Key Vault (premium)
  - Shared device (Luna 7), logical isolation
  - Paas-like integration with Azure services
  - Most affordable, shared hardware, non-confidential
  - Keys and secrets management; no direct PKCS11
- Managed HSM
  - FIPS 140-2 Level 3
  - Dedicated HSM partition (Marvell, isolation at firmware), confidential
  - Paas-like integration with Azure services
  - Keys management; no direct PKCS11 (subset coming for keyless TLS)
- Dedicated HSM
  - Dediated full HSM rack device (Luna 7)
  - No integration with Azure services
  - Intended for full PKCS11 ‘lift and shift’ for on-prem HSMs, general purpose
- Payments HSM (not yet GA – contact keithp@microsoft.com)
  - Intended for credit card payments processors + their regulatory needs
  - Dedicated full HSM rack device
  - No integration for Azure services

Why do customers buy HSMs?

- Vast majority – encryption at rest within Azure services
  - Encrypt my storage account, SQL records, etc
  - Requires deep integration with CSP, high performance requirements
- Keyless TLS
  - Keep your TLS private key in the HSM, popular with customers hosting internet facing servers
- Certificate Authorities / Custom PKCS#11 applications
  - Distribute keys during manufacturing, etc.
- Payments (financial card processing)
  - Strict requirements for performance and security

Why customers buy dedicated

- Lift-and-shift: Need PKC#S11, CAPI, JDA or other protocol due to an app
- Running your own Certificate Authority
- Concerns around access control (build a wall around me)
- Regulatory requirement for physical segregation
- Performance concerns

Why customers don’t like dedicated

- Expensive (especially for high availability, multiple instances)
- Manage it yourself / load balance / failover etc
- Can’t connect it to my PaaS services

Lift-and-shift vs. modernization

- Many customer conversations start with ‘I want to replace my on-prem HSM with one in the cloud’
  - Point existing apps to new code, you are done (not quite)

- Lift-and-shift is possible and a great first step in many cases
  - Does not give the savings that most customers are looking for (engineering time and $$)
  - Often a good way to buy some time between HSMs going out of support and modernizing your apps

- Modernizing apps and rethinking base assumptions usually leads to better outcomes
  - Can you use a shared HSM for this app?
  - You can centralize your control over keys more easily (“crypto center of excellence”)
  - Better options for automating important hygiene tasks (key roll, etc)

HSM as a second-order requirements

- Most customers don't care about HSM – they care about controlling data leaks, preventing the CSP from viewing their data, partitioning data access inside their org, etc
  - HSM is a means to an end
  - Customers are looking well integrated systems that are as cheap as possible
  - Many will ask for dedicated HSMs in this case as an isolation strategy, but in many cases not necessary.
  - Often other solutions can help (data partitioning, better RBAC, data governance, shared HSM, etc)
- Small-medium businesses often are just concerned about checking a compliance box
  - I run my medical practice with Magical Software X, that requires an HSM for HIPPA, what is the cheapest you have?
  - Dedicated is almost never the right answer here

Pitfalls of full ownership/control

- Do you know what all those switches do? ☺
  - Correct usage/maintenance requires a lot of expertise.
  - True experts in HSM/crypto are rare & expensive – adds to TCO
  - True for both on-prem and cloud deployments

- HSM expertise does not always translate into better data management
  - Go for the fancy crypto ciphers etc rather than sensible data governance
  - Increase complexity in your apps
  - No safety net if you crypto-lock yourself out
  - Managed services allow you to rent the device AND the experts

- Analogy – owning your own airplane.
  - Best thing for certain scenarios – but in most cases just buying a ticket is better

Dedicated HSM Challenge: Benefits of ownership without management

- Do you really want *dedicated* or *isolated*?
  - Separate HSM device? Or isolated partition on shared device OK?
  - Separate network?
  - Physical separation for many is a classical view of isolation
  - Can we call anything encrypted by your key as untouchable by me? (Managed HSM philosophy)
- How can I manage your HSM hardware, support infra etc, but without any rights to see/use your keys within?
- How do I prove to you that I cannot see/use the keys within?
  - Regulators will want to see hard proof of the security