White Paper

Security Architecture Models for the Cloud

Introduction

While Hardware Security Module (HSM) customers traditionally have their own infrastructures and data centers and run HSMs on premises, new IT projects are being increasingly implemented in the cloud. This leads to an even higher awareness of data security risks since, in this scenario, the infrastructure runs in third party premises, but the responsibility for data protection remains with the project owner. Strong data protection is becoming ever more important. Cloud-based HSMs in colocation centers can be used for this.

Hardware Security Modules (HSMs) are specialized computing devices that store encryption keys in protected memory. They are also high performance cryptographic “engines” that process data using these keys for high security applications. This paper explains why Utimaco HSMs are agnostic whether they are run in the cloud or on premises and how – technically speaking – service providers can use different “Trust Zones” to maximize the usage of the HSMs independent of the scalability needs of their customers.

HSM deployment scenarios on premises – Single vs. multiple devices

If an HSM is deployed as a single device on a customer’s premises holding vital keys that need to always be available, there are inevitable limitations:

- The HSM may experience a power or hardware failure.
- The demand for key operations may exceed that which a single HSM can deliver.

- Network latency between the HSM and the applications using the keys may be variable or unacceptably high with no guaranteed quality of service.

To overcome these potential limitations, an architecture is required that provides for scalability, but still preserves the security level of a single HSM. This is achieved by the deployment of multiple HSMs, which deliver:

- High availability – active/active or active/standby
- Resilience against downtime by using clusters of HSMs in multiple data centers
- High-speed guaranteed network connectivity and resilient power supplies
- Secure physical protection in data centers with controlled access
- Increased performance with multiple HSMs sharing keys and working together in clusters
- Remote management, which negates the requirement to send staff to the data Center for administrative operations, such as backups
- The option to use an external keystore, where there is no limit on the capacity of the protected key database

High availability – Two different models

An HSM cluster is a number of HSMs all working as one “virtual” HSM in an active/active manner, where the cryptographic load is balanced between all members of the cluster (hence the term “load balancing”). Multiple operational units also provide greater cryptographic processing capacity than just a single unit.

A failover configuration involves one active HSM, and standby or failover units that become operational should the active member become unavailable.

With both models, the internal keystore can be replicated between the HSMs, or a shared external keystore can be used – which of course needs no replication.

Cryptographic control – “Trust Zones”

Cryptographic keys can be separated for management control and for cryptographic access by using the concept of “Trust Zones”. This controls access to specific keys for backup, and for the sharing of keys between cluster members. Control is restricted to those users who have access to the Master Backup Key (MBK), which controls the management of a single global keystore.

“Backup” and “Restore” are important administrative functions that need to be performed whenever a new key is generated. The externally stored backup file is encrypted with the MBK which is held as shares on multiple smartcards. If the HSMs are operated in clusters, the same MBK is shared between the HSMs and the keystore securely replicated. The HSM can be remotely managed using powerful, but simple, utilities that enable standard administrative functions such as logging, upgrading firmware, backup, adding users, etc.

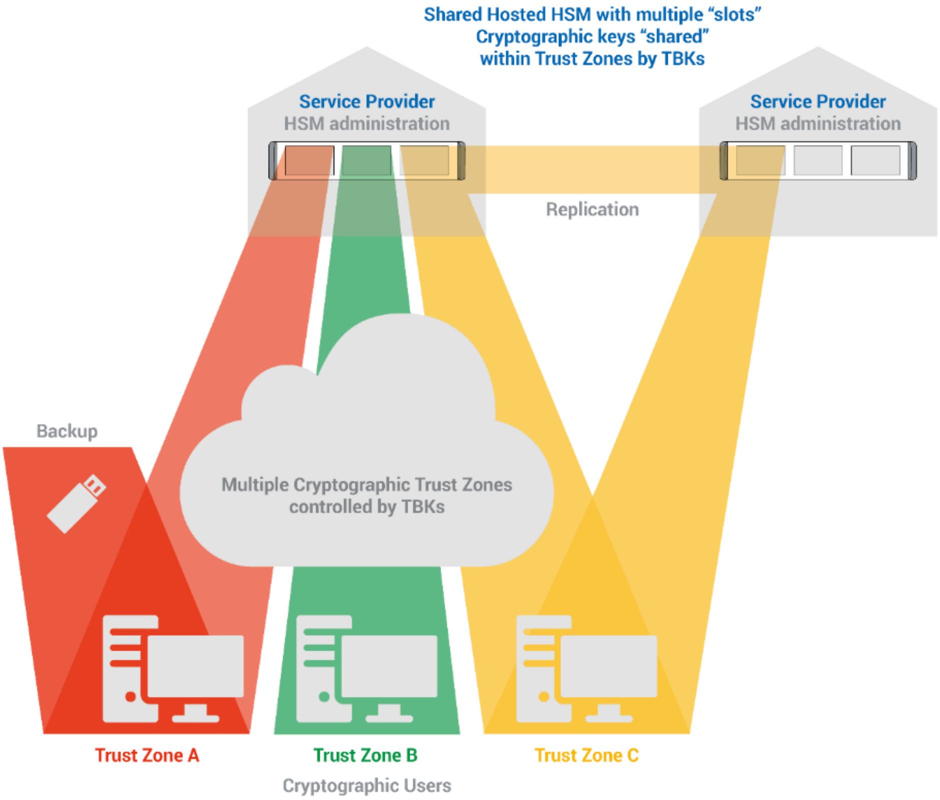

HSM deployment scenarios – On premises vs. in the cloud

Whilst some organizations have their own infrastructure, data centers and IT staff resources, many organizations find it difficult to provide all these facilities. These requirements therefore lead to the adoption of the "cloud" model of HSM deployment, where service providers host clusters of HSMs in multiple data centers. These are provided "as a Service" to cryptographic users/customers – organizations that need HSMs to protect their keys. The data centers may be operated by the service provider or by another third party. The customer's application, whether deployed on their own premises or hosted by a cloud application service provider, will communicate with the HSM using an encrypted IP connection, with no changes needed to the application.

It should be noted that the HSM operates technically in the same way whether it is owned by the cryptographic user or whether it is provided "as a Service" – only the way in which it is managed is different.

Clearly, using a service provider to hold and manage cryptographic keys in an HSM hosted by a third party involves a considerable degree of trust on the part of the customer. As a result, a clear separation of duties and responsibilities has to be agreed and documented between the service provider and the customer to ensure secure operation. The security provided by an HSM comes as much from the policies and procedures under which it is managed as from the HSM's hardware and software. Such policies include defining user roles for different kinds of functionalities.

Separation of duties should include multiple people authenticating to enable a particular role. This separation of duties is broken down into two main groups of users:

- HSM administrators
- Cryptographic key managers

The separation of duties between the groups of HSM users/administrators and the cryptographic users/managers is enabled by the issuing of secure tokens to the people holding these respective roles. Different roles permit different functionalities to be performed. Multiple token holders may be required to grant permission for a role using the “M of N” authentication method.

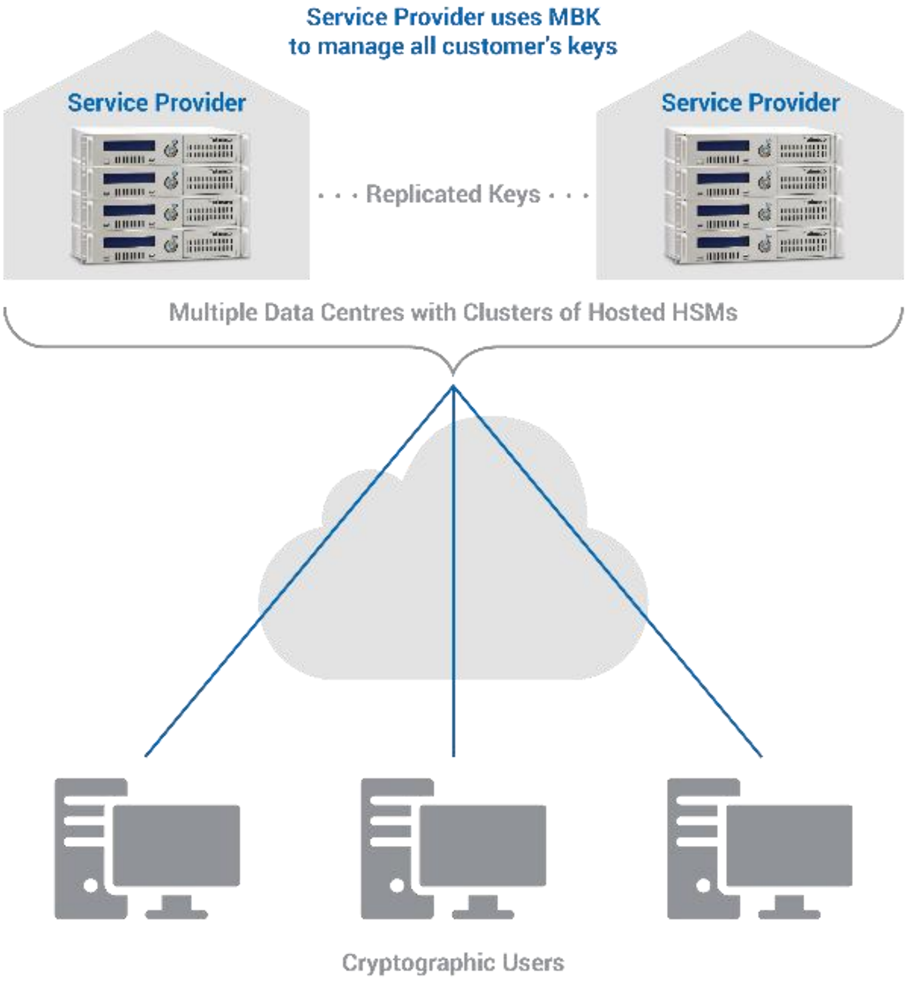
Service Provider uses MBK
to manage all customer's keys
Service Provider
Service Provider
· · · Replicated Keys · · ·
Multiple Data Centres with Clusters of Hosted HSMs
Cryptographic Users

HSM administration – Service provider or customer in the lead

However, cryptographic management of a key is determined by the key's Trust Zone – it controls to where cryptographic keys can be backed up, or to where they can be shared for clustering. Control is managed by who holds the Trust Zone backup key.

Therefore, an important decision has to be made by the customer as to whether they want the service provider to control the Trust Zone, or whether they want to retain this control for themselves.

Cloud "Exclusive" – Fully managed model

If the customer has exclusive use of an HSM, they may want the service provider to be responsible for key management, in which case the service provider will manage the MBK to backup the HSM's global key database. The customer may, however, wish to perform their own backups, and they would then hold the MBK. In both cases, there will be one Trust Zone controlled by one MBK under the management of either the service provider or the customer.

Cloud "Shared" – Customer key management model

To maximize the use of an HSM, the service provider may decide to offer it as a "shared" service by dividing the keystore into cryptographically separate "slots" or "partitions" that each operate like mini-HSMs. In this case, backing up and sharing keys between cluster members is managed by the Tenant Backup Key (TBK) that controls the access to each slot. In this case, there will be multiple Trust Zones controlled by multiple TBKs that can either be under the management of the service provider or the customer.

Utimaco HSMs in the Cloud

The use of an HSM from Utimaco as a cloud service allows for all advantages provided by Utimaco, namely:

- The flexible management of user groups (roles) and the methods of authentication.
- The HSM's functionality being customizable using the Software Development Kit (SDK) to build firmware modules to provide specialized functions, such as non-standard algorithms.

About Utimaco

Utimaco is a leading manufacturer of HSMs that provide the Root of Trust to all industries, from financial services and payment to the automotive industry, cloud services to the public sector. We keep cryptographic keys and digital identities safe to protect critical digital infrastructures and high value data assets. Our products enable innovations and support the creation of new business by helping to secure critical business data and transactions.

Founded in 1983, Utimaco HSMs today are deployed across more than 80 countries in more than 1,000 installations. Utimaco employs a total of 200 people, with sales offices in Germany, the US, the UK and Singapore. For more information, visit https://hsm.utimaco.com

Would you like to try out our HSM or implement new algorithms yourself?

Download the Utimaco Software HSM Simulator to get started immediately learning about HSM devices. It is a FREE fully functioning Software version of the Hardware HSM. The download package includes documentation on our product. Register here for your download.

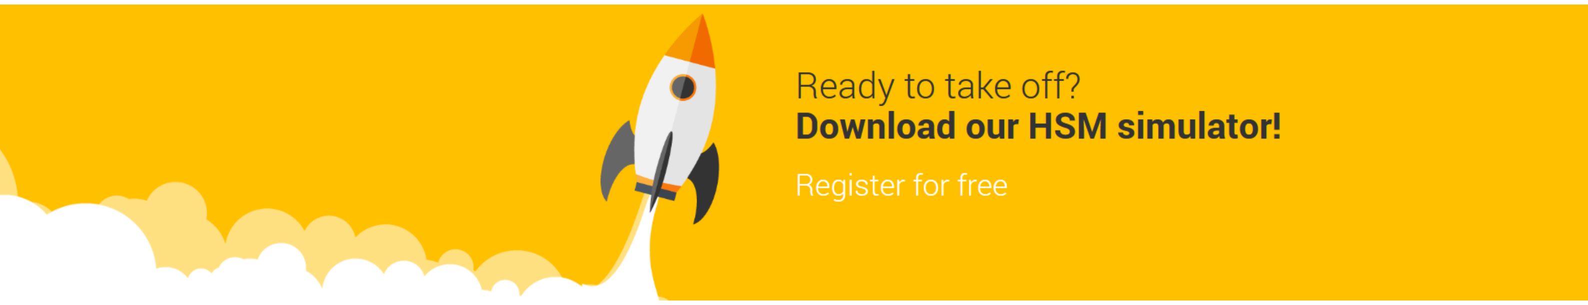

The Simulator download includes tools for creating user accounts, sample code and libraries for PKCS#11 Microsoft CNG, Java JCE and the Utimaco CXI API to link and test your code. The Simulator will run on a Windows or Linux platform. This is your opportunity to try out the Utimaco HSM technology easily and without initial cost. Register for free: https://hsm.utimaco.com/downloads/utimaco-portal/hsm-simulator