LAYER SEVEN SECURITY

SECURITY IN SAP HANA

THE CHALLENGES OF IN-MEMORY COMPUTING

WHITE PAPER

SECURITY IN SAP HANA

THE CHALLENGES OF IN-MEMORY COMPUTING

CONTENTS

INTRODUCTION

According to research performed by the International Data Corporation (IDC), the volume of digital information in the world is doubling every two years. The digital universe is projected to reach 40,000 exabytes by 2020. This equates to 40 trillion gigabytes or 5200 gigabytes for every human being in the world in 2020. [1] As much as 33 percent of this information is expected to contain analytic value. Presently, only half of one percent of available data is analyzed by organisations.

The extraction of business intelligence from the growing digital universe requires a new generation of technologies capable of analysing large volumes of data in a rapid and economic way. Conventional approaches rely upon clusters of databases that separate transactional and analytical processing and interact with records stored in secondary or persistent memory formats such as hard disks. Although such formats are non-volatile they create a relatively high level of latency since CPUs lose considerable amounts of time during I/O operations waiting for data from remote mechanical drives. Contemporary persistent databases use complex compression algorithms to maximise data in primary or working memory and reduce latency. Nonetheless, latency times can still range from several minutes to days in high-volume environments. Therefore, persistent databases fail to deliver the real-time analysis on big data demanded by organisations that are experiencing a significant growth in data, a rapidly changing competitive landscape or both.

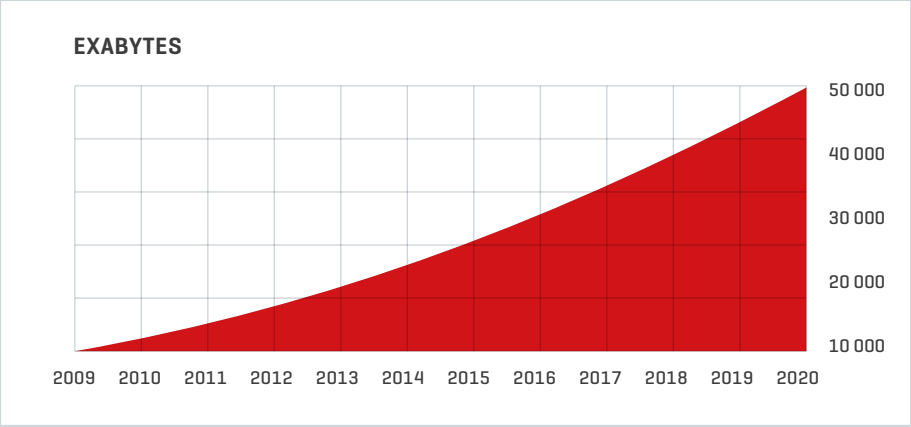

Figure 1.1: Worldwide Digital Information

In-memory databases promise the technological breakthrough to meet the demand for real-time analytics at reduced cost. They leverage faster primary memory formats such as flash and Random Access Memory (RAM) to deliver far superior performance. Primary memory can be read up to 10,000 times faster than secondary memory and generate near-zero latency. While in-memory technology is far from new, it has been made more accessible by the decline in memory prices, the widespread use of multi-core processors and 64-bit operating systems, and software innovations in database management systems.

[1] The Digital Universe in 2020, IDC, 2012

PREVENTION	DETECTION
Authentication and Authorization	Data Discovery and Classification
Database Firewall	Privilege Analysis
Encryption	Configuration Management
Data Redaction and Masking	Logging and Auditing
Patch Management	

Figure 1.2: Database Security

The SAP HANA platform includes a database system that processes both OLAP and OLTP transactions completely in-memory. According to performance tests performed by SAP on a 100 TB data set compressed to 3.78 TB in a 16-node cluster of IBM X5 servers with 8 TB of combined RAM, response times vary from a fraction of a second for simple queries to almost 4 seconds for complex queries that span the entire data range. [2] Such performance underlies the appeal and success of SAP HANA. Since its launch in 2010, SAP HANA has been become SAP's fastest growing product release.

SAP HANA has emerged against a backdrop of rising concern over information security resulting from a series of successful, targeted and well-publicized data breaches. This anxiety has made information security a focal point for business leaders across all industry sectors. Databases are the vessels of business information and therefore, the most important component of the technology stack. Database security represents the last line of defense for enterprise data. It should comprise of a range of interdependent controls across the dual domains of prevention and detection. Illustrative controls for each domain are provided in Figure 1.2.

The most advanced persistent databases are the product of almost thirty years of product evolution. As a result, today's persistent databases include the complete suite of controls across both domains to present organisations with a high degree of protection against internal and external threats. In-memory databases are in comparison a nascent technology. Therefore, most do not as yet deliver the range of security countermeasures provided by conventional databases. This includes:

- Label based access control;
- Data redaction capabilities to protect the display of sensitive data at the application level;
- Utilities to apply patches without shutting down databases;
- Policy management tools to detect database vulnerabilities or misconfigurations against generally-accepted security standards.

The performance edge enjoyed by in-memory database solutions should be weighed against the security disadvantages vis-à-vis persistent database systems. However, it should be noted that the disadvantages may be short-lived. Security in in-memory databases has advanced significantly over a relatively short period of time. For example, SPS 06 introduced a number of security enhancements not available in SPS 05, released only seven months earlier. This includes support for a wider number of authentication schemes, the binding of internal IP addresses and ports to the localhost interface, a secure store for credentials required for outbound connections and more granular access control for database users.

[2] SAP HANA Performance, SAP AG, 2012

The most crucial challenge to database security presented by the introduction of in-memory databases is not the absence of specific security features but architectural concerns. Server separation is a fundamental principle of information security enshrined in most control frameworks including, most notably, the Payment Card Industry Data Security Standard (PCI DSS). [3] According to this principle, servers must be single purpose and therefore must not perform competing functions such as application and database services. Such functions should be performed by separate physical or virtual machines located in independent network zones due to differing security classifications that require unique host-level configuration settings for each component. This architecture also supports layered defense strategies designed to forestall intrusion attempts by increasing the number of obstacles between attackers and their targets. Implementation scenarios that include the use of in-memory databases such as SAP HANA as the technical infrastructure for native applications challenge the principle of server separation. In contrast to the conventional 3-tier architecture, this scenario involves leveraging application and Web servers built directly into SAP HANA XS (Extended Application Services). Unfortunately, there is no simple solution to the issue of server separation since the optimum levels of performance delivered by in-memory databases rely upon the sharing of hardware resources between application and database components.

Aside from such architectural concerns, the storage of large quantities of data in volatile memory may amplify the impact of RAM-based attacks. Although widely regarded as one of the most dangerous security threats, attacks such as RAM-scrapping are relatively rare but are becoming more prevalent since attackers are increasingly targeting volatile memory to circumvent encrypted data in persistent memory. Another reason that RAM-based attacks are growing in popularity is that they leave virtually no footprint and are therefore extremely difficult to detect. This relative anonymity makes RAM-based attacks the preferred weapon of advanced attackers motivated by commercial or international espionage.

[3] PCI DSS 2.0 (2010), Requirement 2.2.1: Implement only one primary function per server to prevent functions that require different security levels from co-existing on the same server

NETWORK SECURITY

SAP HANA should be located in a secure network zone with minimal connections to other zones. Network connectivity should be limited to the services required for each implementation scenario. Table 2.1 lists the most common internal and external connections. Note that xx represents the SAP HANA instance number.

External inbound connections include the SQLDBC protocol for database clients and data provisioning (3xx15, 3xx17), and administrative functions performed through the SAP HANA Studio (5xx13, 5xx14, 1128, 1129). They also include HTTP and HTTPS for Web-based access to SAP HANA XS and other components (80xx, 43xx). Connections to the hdbrss binary through SAProuter (3xx09) are deactivated by default and should only be enabled in specific support cases.

External outbound connections should be limited to the SAP Solution Manager from the diagnostics agent installed on each system, the SAP Service Marketplace from the SAP HANA Lifecycle Manager, and required calls to external servers from SAP HANA XS. Connections for smart data access and integration for R environments should only be enabled when required.

Internal communications between components within a single host system or multiple hosts in a distributed system should be performed within a dedicated private network using separate IP addresses and ports that are isolated from the rest of the network. To this end, the default setting that blocks access from external network hosts by binding internal IP addresses and ports to the localhost interface in single-host scenarios should not be modified. The default port for the localhost is 3xx00. The default port range for hosts in a distributed system is 3xx01-3xx07.

Internal communications should be limited to links between components within the same host, recognised hosts in a distributed environment, and primary and secondary sites for replication purposes. VPN or IPSec can be used to secure the communication channel between primary and secondary sites.

SOURCE	DESTINATION
SAP Solution Manager Diagnostics Agent (SMD)	SAP Solution Manager
SAP HANA Lifecycle Manager	SAP Service Marketplace
SAP HANA XS	External Servers
SAP Smart Data Access	External data sources
SAP HANA	R environments

Figure 2.1: Outbound Connections

PROTOCOL	TCP PORT	CLIENTS
SQLDBC (ODBC/JDBC)	3xx15 3xx17 3xx13 3xx14 1128 1129	Application servers SAP HANA Studio End users Replication systems
HTTP[S]	80xx 43xx	Web browsers Mobile devices SAP HANA Direct Extractor Connection (DXC)
Internal / Proprietary	3xx09	SAP Support

Figure 2.2: Inbound Connections

AUTHENTICATION AND AUTHORIZATION

SAP HANA supports a wide range of authentication methods. The most basic is username/ password combinations for database users created and maintained through the SAP HANA Studio, command line interfaces such as hdbsql or through NetWeaver Identity Management (IdM). User data is stored in a local repository.

External user repositories such as Kerberos and Security Assertion Markup Language (SAML) can be used to authenticate access to SAP HANA through database clients such as SAP HANA Studio or front-end applications such as Business Intelligence, Business-Objects and CRM. However, external repositories still require database users since user identities are mapped to identities in SAP HANA.

Client certificates issued by a trusted Certification Authority (CA) can be used to authenticate users accessing the database through SAP HANA XS using HTTP. SAP logon tickets issued by SAP systems such as the NetWeaver Application Server or Portal can also be used to authenticateWeb-based access through SAP HANA XS. Both options require the configuration of the Secure Sockets Layer (SSL) protocol.

Kerberos and SAML are generally more secure authentication schemes than client certificates and logon tickets. However, all four methods can be used for Single Sign-On (SSO). SAP HANA XS includes tools for configuring and maintaining authentication schemes. Kerberos requires the installation of client libraries within the SAP HANA host system and mapping of database users to external identities in the Kerberos key distribution center (KDC).[4]

The use of SAML for user authentication involves configuring identity providers and mapping external and database users using SAP HANA XS. Hash or signature algorithms such as SHA-1, MD5 and RSA-SHA1 or X509Certificate elements should be used to secure XML signatures used in SAML assertions and responses.

Direct authentication in SAP HANA requires the configuration of a strong password policy maintained in the indexserver.ini system properties file. This file should be maintained through the SAP HANA studio. Therefore, direct changes to the file should be avoided. A major drawback of password policies in SAP HANA is that changes to the indexserver.ini file cannot be audited. Table 3.1 outlines password parameters, default configurations and recommended settings. Password policies can be reviewed through the M_PASSWORD_POLICY system view. Note that the maximum_invalid_connect_attempts parameter does not apply to SYSTEM users. Therefore, such users are more likely to be targeted for brute-force or other password attacks. Technical users can be excluded from the maximum_password_lifetime check through the SQL statement ALTER USER <user_name> DISABLE PASSWORD LIFETIME.

Forbidden passwords should be specified in the _SYS_PASSWORD_BLACKLIST (_SYS_SECURITY) table. The table is empty by default. SAP HANA supports blacklisting of passwords based on either exact matches or keywords contained within passwords. Blacklisted words can be either case-sensitive or case-insensitive.

[4] Refer to SAP Note 1837331

PARAMETER	DEFAULT VALUE	RECOMMENDED VALUE
minimal_password_length	8	8
password_layout	Aa1	A1a_
force_first_password_change	true	true
last_used_passwords	5	6
maximum_invalid_connect_attempts	6	4
password_lock_time	1440	1440
minimum_password_lifetime	1	1
maximum_password_lifetime	182	90
maximum_unused_initial_password_lifetime	28	5
maximum_unused_productive_password_lifetime	365	30
password_expire_warning_time	14	14

Figure 3.1 – Password Parameters

Direct assignment of authorizations to database users should not be performed. Rather, permissions should be granted through predefined roles. SAP HANA includes several standard roles designed to meet most business scenarios and provide a template for custom role development. Users require both the privileges to perform a specific action and access to the relevant object to perform database operations. Privileges are categorized into several classes. System privileges are equivalent to SQL permissions for administrative tasks including schema creation, user management and backup and recovery. Object privileges are used to control actions such as SELECT, CREATE, ALTER etc. at the object level. Analytic privileges are used to enforce context-dependant access to data in information models. This ensures that database users are only able to access database objects for their specific company, region or other variables. The _SYS_BI_CP_ALL privilege can override other analytic privileges when combined with the SELECT object privilege. This combination can give users access to all data in every data set. Therefore, SAP does not recommend the use of _SYS_BI_CP_ALL, especially in production systems.

Standard users delivered with SAP HANA should be secured after the initial install or upgrade. The password provided by the hardware vendor for the <sid>adm operating system user should be changed after the handover. A password reset should also be performed for the powerful SYSTEM user which should then be deactivated. This user should not be employed for administrative operations post install or upgrade.

In multi-tenant databases, the default DB-level isolation should be supported by OS-level measures to prevent cross-database attacks. This includes using separate OS users to ensure database processes for multiple tenants do not run with the identical <sid>adm user.

ENCRYPTION

The TLS/ SSL protocol should be used to encrypt client-server traffic and internal communications in SAP HANA. SSL is not invulnerable. SSL proxies are widely available and can be used to intercept and decrypt packets passed between endpoints within a network. Despite these and other limitations, SSL remains the most common method for cryptographically securing network communications.

Implementing SSL for client-server SQL traffic in SAP HANA requires both client and server side configuration. The OpenSSL library or the SAP Cryptographic Library can be used to create the required public-key certificates. However, SAP recommends the former which is installed by default. Public and private key pairs and corresponding certificates are stored in the personal security environment (PSE) within each server. SSL parameters are maintained in the indexserver.ini configuration file. The sslCreate-SelfSignedCertificate parameter should be set to false to prevent the use of self-signed certificates.

The use of SSL for internal communications between hosts in a distributed environment involves configuring server and clients PSEs in each host. A reputed Certification Authority should be used to sign certificates used for internal communications.

Separate TLS/ SSL certificates and keys should be configured for each database in multi-tenant environments. Shared trust stores with wildcard server certificates could enable users from one database to log on to other databases.

The SAP Web Dispatcher should be configured to support HTTPS (HTTP over SSL/TLS). HTTPS requests should either be re-encrypted before they are forwarded to the ICM within each NetWeaver Application Server instance or forwardedwithout unpacking for end-to-end SSL. SSL termination isthe least preferred option. Therefore, the wdisp/ssl_encrypt option in the icm/server_port_<xx> parameter should be set to 1, 2 or ROUTER. This requires installation of SSL libraries and following the detailed configuration procedures provided by SAP. Once SSL is implemented, SAP HANA XS should be configured to refuse non-HTTPS connection requests through the sslenforce option in the runtime configuration.

In-memory data in SAP HANA is automatically replicated to an internal persistent data volume for recovery purposes. Data volumes can be encrypted with a 256-bit strength AES algorithm. To enable persistence encryption in existing installations, SAP recommends uninstalling and reinstalling the database. The alternative method involving generating root encryption keys using the hdbnsutil program without reinstalling the database may not encrypt all data. In either case, the SQL command to enable encryption is ALTERSYSTEM PERSISTENCE ENCRYPTION ON. The activation of persistence encryption can be verified in the ENCRYPTION_ACTIVE_AFTER_NEXT_SAVEPOINT column which should contain the value TRUE.

Root encryption keys are stored using the SAP NetWeaver secure storage file system (SSFS). Although SAP HANA generates new and unique root keys during installation, keys should be changed using the command line tool rsecssfx immediately after the handover from hardware partners to ensure they are not shared with external parties.

Data volume encryption keys should be periodically changed using hdbnsutil. For further information, refer to Notes 2183624 and 2228171

Passwords for database users are obfuscated with the SHA-256 hash function. However, database redo log files containing the history of changes made to the database are not encrypted in persistent volumes.

Other than data in the secure internal credential store, database backups are also not encrypted. The same applies to database traces. Therefore, SAP HANA installations containing sensitive data should be supplemented with third-party solutions for the encryption of files at the operating system level and data backups. Furthermore, the use of tracing functions should be minimized and limited to short-term analysis. Trace files can be identified through the file extension .trc and can be deleted using the Diagnosis Files tab in the SAP HANA Administration editor. The number and size of trace files can be restricted by adjusting the maxfiles and maxfilesize parameters for trace file rotation in the global.ini file for all services or the indexserver.ini file for individual services.

AUDITING AND LOGGING

<Event Timestamp>
<Service Name>
<Hostname>
<SID>
<Instance Number>
<Port Number>
<Client IP Address>
<Client Name>
<Client Process ID>
<Client Port Number>
<Audit Level>
<Audit Action>
<Active User>
<Target Schema>
<Target Object>
<Privilege Name>
<Grantable>
<Role Name>
<Target Principal>
<Action Status>
<Component>
<Section>
<Parameter>
<Old Value>
<New Value>
<Comment>
<Executed Statement>
<Session Id>

Figure 5.1 – Fields in Audit Entries

Enabling auditing in SAP HANA requires the AUDIT ADMIN or INFILE ADMIN system privilege and should be performed either through Systems Settings for Auditing or the SQL statement ALTER SYSTEM LOGGING ON. The global_auditing_state parameter in the global.ini file will display the value true if logging has been successfully enabled.

Once enabled, audit policies should be configured to log actions that include SELECT, INSERT, UPDATE, DELETE, EXECUTE and other statements when combined with specific conditions. Policies can be configured for specific users, tables, views and procedures. It is recommend to audit all actions performed by privileged users including the SYSTEM user and actions that impact sensitive database objects.

Polices are created and maintained in the Audit Policies area in the Auditing tab of the Security editor. Polices can be configured to log all SQL statements or successful/unsuccessful attempts. A severity level must be assigned for each audit policy. The options include EMERGENCY, ALERT, CRITICAL, WARNING and INFO. These ratings are important triggers that impact the communication, escalation and resolution of audit events.

The fields captured in audit entries are detailed in Figure 5.1. Note that only actions performed directly through HANA are logged. Therefore, actions such as configuration changes performed through the OS layer will not be recorded in the HANA log files.

The parameter default_audit_trail_type should not be set to CSVTEXTFILE in productive systems. Unless the default_audit_trail_path is modified, audit entries will be written to the same directory as trace files (/usr/sap/<sid>/<instance>/<host>/trace). Therefore, log entries can be read by database users with DATA ADMIN, CATALOG READ, TRACE ADMIN or INFILE ADMIN privileges, as well as operating system users in the SAPSYS group. The value CSTABLE or SYSLOGPROTOCOL should be used for the default_audit_trail_type parameter. This will ensure that SAP HANA uses internal database tables or the operating system syslog for the storage of the audit trail.

The syslog daemon should be configured to log entries in a central server or receiver in distributed environments. The max_log_file and max_log_file_action parameters in the /etc/sysconfig/auditd file should be used to configure an appropriate file size and rotate logs to ensure uninterrupted service.

The syslog protocol can be used to support the secure storage of audit logs from SAP HANA by preventing database administrators from accessing and modifying log files. It also provides a widely-recognized format for event analysis and reporting and therefore provides for seamless integration with a variety of open source and commercial security information and event management (SIEM) systems. However, syslog should be used with IPSEC or SSH port tunnelling to secure log transmissions and protect the integrity of log data.

SAP HANA APPLIANCE

SAP HANA is delivered as an appliance. The infrastructure includes a fully installed and configured SUSE Linux Enterprise (SLES) or Red Hat Enterprise Linux (RHEL) operating system. Although security settings are pre-configured by SAP, they should be verified against the security and hardening guidelines issued by SUSE or Red Hat. Note that changes to the underlying platform of SAP HANA may impact terms of support and therefore should be reviewed and agreed beforehand with SAP and hardware partners.

Security settings in SUSE Linux distributions are managed through the Security Center and Hardening module of the setup and configuration tool YaST. There are several predefined security configurations available in the module. The Network Server option provides the most secure default configuration for application and database servers.

Password rules should be configured in line with organisational requirements for age, complexity, length and other variables. The option for dictionary and noun checks for new passwords should be selected. The default password encryption method is the Blowfish cipher. Alternative methods such as AES, Twofish and Threefish are less susceptible to attack. The display of the login prompt should be delayed by 1-2 seconds following an unsuccessful login attempt. Furthermore, the options for logging unsuccessful login attempts should be enabled and remote access to the graphical login manager should be disabled. The secure filepermission setting is recommended for networked systems.

Kernel parameters in the /etc/sysctl.conf file should be configured to secure Linux against common attacks. This includes enabling TCP SYN cookie protection, IP spoofing protection and virtual address space randomization, and disabling IP source routing.

The netstat service can be used to review open ports and services. Vulnerable services such as ftp, telnet and sendmail should be deactivated. OpenSSH and postfix can be used as secure alternatives for such services. For SSH, direct root logins should be disabled, privilege separation enabled and only version 2 of the protocol should be accepted. Unnecessary software packages (RPMs) should also be identified and removed. This will minimize the attack surface and streamline maintenance.

The xinetd program can be deployed as a TCP wrapper to filter incoming requests as a substitute for network firewalls. This is performed through access control lists configured in the /etc/hosts.allow and /etc/hosts.deny files. The former takes precedence over the latter. Therefore, the /etc/hosts.deny should be configured with a deny-all rule and requests from specific hosts should be allowed in the /etc/hosts.allow file. Xinetd provides the flexibility to regulate connections based on hostname, IP address, subnet, service and other variables. Note that both the aforementioned files are world readable. Therefore, the file permissions should be modified. Permissions should also be changed for world-writable files and directories which can be modified by all users. This excludes files in /tmp and other directories that can only be changed by file owners.

Access to the root password and therefore, the ability to execute root commands should be tightly controlled. System administrators that require the ability to perform root commands should only be provided root access through sudo. SUSE Linux Enterprise enables organisations to specify allowable root commands for sudo users. This is managed through the /etc/sudoers file.

Although Yast can be used to encrypt partitions during installation, it is not recommended to encrypt running systems since this will erase all data in target partitions. Therefore, encrypted container files should be used to protect files and folders with sensitive data. This can be performed through YaST Expert Partitioner – Crypt Files+Add Crypt File – Create Loop File.

SUSE includes a system integrity analyzer called Seccheck. This executes daily, weekly and monthly shellscripts to review security-related configuration settings. The dailyscript reviews password parameters, root users, aliases and .rhosts. The weekly and monthly scripts perform more exhaustive checks but require the installation of a password cracking tool on the host that may present a security risk.

SUSE also includes AIDE for file and system integrity monitoring. The use of AIDE is highly recommended but is not a default component of the SUSE installation. The AIDE binary that is used as the configuration database for integrity checks should be installed on a separate host. The database is installed in the /var/lib/aide/ directory and attributes are defined in the /etc/aide.conf configuration file. Aide checks are performed through the command aide –check. The parameter –v should be run to display the detail of the differences identified between the database and file system.

SUSE security patches should be retrieved through the Novell Subscription Management Tool (SMT) rather than YaST Online Update. SMT should distribute patches to clients from a proxy system to avoid direct outbound connections to the Novell Customer Center and nu.novell.com servers. A proxy system should also be used to synchronize time across all servers. A dedicated NTP server should obtain time from at least two external authoritative sources which should then be distributed to internal servers as part of a defined time synchronization hierarchy.

Hardening for Red Hat Enterprise Linux Enterprise platforms should follow a similar approach as SUSE Linux distributions. This includes removing services such as telnet, rlogin, NFS and SMB. Send packet redirects, source routed packet acceptance and ICMP redirect acceptance should be disabled. On the other hand, Bad Error Message Protection, Ignore Broadcast Requests and TCP/SYN cookies should be enabled. The password hashing algorithm should be set to SHA-512 or higher, rather than the insecure MD5. Strong password policies should be defined using password parameters in the Pluggable Authentication Modules (PAM). Access to PAM configuration files in the /etc/pam.d/* directory should be restricted. Finally, SELinux should be configured to create confined domains for daemons. This would limit the impact of buffer overflow and other attacks that target background processes.

SAP HANA ONE

SAP HANA One is an Infrastructure as a Service (IaaS) solution that uses the Amazon Web Services (AWS) public cloud for the hosting of SAP HANA. Since IaaS provides organisations with the ability to partially control the configuration of cloud infrastructure including networking components, operating systems, storage and deployed applications, many of the recommendations provided in earlier sections of this paper can be applied to SAP HANA One.

However, there are certain challenges unique to cloud services that deserve specific consideration. Similar to all cloud-based solutions, SAP HANA One provides customers with the opportunity to rapidly deploy and scale services while minimizing capital and operational costs by leveraging the economies of scale provided by shared IT resources. However, security concerns are an inhibiting factor that have contributed to the relatively low rate of adoption of cloud computing.

These concerns include cross-border data flows. The AWS global infrastructure spans a variety of geographic zones including countries in Asia, Europe and North and South America. Therefore, data flows may not be contained within the customer's country of origin. This may expose organisations to country or region-specific regulations governing privacy and other areas, as well as international litigation in the event of information leakage. Therefore, contractual agreements for cloud services should be closely reviewed by legal representatives and include assurances that electronic data and copies of data are stored in specific geographic locations. Agreements should also include a right-to-audit clause or terms for the regular provision of evidence of compliance against specific information security requirements.

Another concern is virtualization, which is a key enabler of economies of scale in multitenant cloud services. Virtualized operating systems such as virtual machines (VM) must be compartmentalized, isolated and hardened. Furthermore, since network firewalls are incapable of inspecting communications between VMs on the same host, virtual firewalls should be deployed at the hypervisor level. Alternatively, customers can physically isolate hardware by using dedicated instances in the AWS cloud.

The EC2-VPC (Virtual Private Cloud) offers a more secure deployment option for SAP HANA One than EC2-Classic through logical separation within a virtual network, enabling complete control of IP address ranges, subnets, routing tables and network gateways. Customers are therefore able to control inbound and outbound connections to subnets using ACLs. Network Address Translation (NAT) should be used to create a private subnet and prevent direct access to SAP HANA One from the Internet. The bridge between cloud infrastructure and onsite datacenters should be secured through VPN. Finally, firewall policies applied through AWS security groups should be configured to only enable permitted clients to access the required ports and services.

CONCLUSION

SAP HANA represents a significant technological breakthrough by delivering real-time analytics for large volumes of data more effectively and practicably than traditional persistent databases. This performance can be leveraged to provide new insight, drive rapid innovation and meet the strategic challenges of the growing digital universe. However, realizing the greatest performance benefits of SAP HANA requires an architecture that challenges one of the fundamental principles of information security. Furthermore, the reliance upon primary memory for both analytical and transactional processing may expose in-memory systems to complex, powerful and difficult to detect attacks. Finally, since in-memory computing is an emerging technology, SAP HANA does not currently provide the wide array of security capabilities available in the most advanced persistent databases.

Notwithstanding these concerns, the secure deployment and administration of SAP HANA requires the application of multiple, interdependent controls in the areas of network and communication security, authentication and authorization, data encryption and auditing and logging. It also requires the secure configuration of the Linux platform. Cloud-based scenarios require measures to control cross-border data flows, virtualized operating systems and virtual private clouds. A comprehensive security framework comprising these areas will support the safeguarding of information assets in the SAP HANA in-memory database.

LAYER SEVEN SECURITY

Layer Seven Security empowers organisations to realize the potential of SAP systems. We serve customers worldwide to secure systems from cyber threats. We take an integrated approach to build layered controls for defense in depth.

CONTACT US

Westbury Corporate Centre
2275 Upper Middle Road East, Suite 101
Oakville, Ontario, L6H 0C3, Canada
Tel. (Toll Free): 1 888 995 0993
Tel. (Office): 905 491 6950
Fax.: 905 491 6801
E-mail: info@layersevensecurity.com
www.layersevensecurity.com

SAP Partner