

JOURNEY TO CLOUD

Volume 2, Issue 1

The Next Generation

- BIG DATA, SMALL BUDGET
- HADOOP
- COMPATIBLEONE* NEXT-GENERATION CLOUD MANAGEMENT
- CLOUD DEPLOYMENT AND DESIGN ON INTEL PLATFORMS
- EXTREME FACTORY* HPC ON-DEMAND SOLUTION
- DISTRIBUTED HEALTH CLOUD ARCHITECTURE
- DEFINING AN EFFECTIVE CLOUD COMPUTING
STORAGE STRATEGY
- CLOUD SECURITY





CLOUD COMPUTING
has moved from
HYPE to **KEY IT**
STRATEGY, offering
SAVINGS,
SCALABILITY, and
FLEXIBILITY for
enterprises of all
types and sizes.

CONTENTS

BIG DATA, SMALL BUDGET

Handling the Explosion of Unstructured Data with a Scale-Out Storage Architecture 1

HADOOP

A Distributed, Parallelized Platform for Handling Large Data 7

COMPATIBLEONE

Next-Generation Cloud Management 15

CLOUD DEPLOYMENT AND DESIGN

On Intel Platforms 25

EXTREME FACTORY

In an Unclear HPC Cloud Landscape, an Efficient HPC-on-Demand Solution 35

DISTRIBUTED HEALTH

Cloud Architecture 41

DEFINING AN EFFECTIVE CLOUD COMPUTING STORAGE STRATEGY 47

CLOUD SECURITY

Securing the Infrastructure with Intel® Trusted Execution Technology 53

The Age of Data-Intensive Computing

Now more than ever, companies are faced with big data streams and repositories. Aggregating, reading, and analyzing large amounts of structured, multi-structured, and complex or social data is a big challenge for most enterprises. According to *The Guardian*, over the next decade there will be 44 times more data and content than today. Information is exploding in volume, variety, and velocity. But the determining factor is the fourth V, the ability to extract value from available data. The large volume of data is accompanied by a diversification in types of information and the speed and frequency with which it's being generated and distributed. Demand for real-time and near-real-time data processing has increased significantly. And that demand alone presents a real challenge to an already overstretched IT infrastructure.

Because of these challenges, there's a need for an alternative approach to data processing and business intelligence. This alternative approach includes both the Hadoop Map Reduce* framework and unstructured data warehousing solutions. This new approach doesn't invalidate the traditional data warehouse, but it does acknowledge its limitations in dealing with large volumes of data.

In this issue, we explore alternative solutions for big data scale-out storage solutions, Hadoop, next-generation cloud management, cloud security, HPC on demand, and more. Be sure to let us know what you think.

Parviz Peiravi
Editor in Chief
parviz.peiravi@intel.com

Sally Sams
Production Editor
sally.sams@intel.com



BIG DATA, SMALL BUDGET

Suhas Nayak
Storage Solutions Architect
Intel Corporation
suhas.nayak@intel.com

Chinmay S. Patel
Cloud Marketing Manager
Intel Corporation
chinmay.s.patel@intel.com

Billions of connected users, and billions of connected devices, are generating an unprecedented amount of data in the form of digital images, HD videos, music, sensory data, log files, and emails, among others.

STORAGE REQUIREMENTS

This explosion of data, classified as unstructured data (or big data)—along with longer data retention requirements, stringent service level agreements (SLAs) and high utility costs—puts significant demand on IT infrastructure.

Without technological innovation, this would translate into significantly higher capital and operational expenses for IT organizations still coping with the effects of unstable global economies.

According to IDC, our digital universe will grow to be 2.7 zettabytes by end of 2012, since the amount of data is doubling every two years.¹ At this rate, there will be more than 35 zettabytes of data by end of 2020, about 90 percent of which will be unstructured data.


While data is growing rapidly, IT budgets are relatively shrinking and IT managers are asked to do more with less. The industry is embracing

the explosion of unstructured data with a new type of storage architecture called scale-out storage architecture.

As shown in Figure 1, traditional enterprise storage is implemented using proprietary solutions in a centralized storage area network (SAN). This type of infrastructure, often called scale-up storage architecture, has limited scalability and higher costs in the context of the growth in unstructured data.

While this type of enterprise storage is quite useful for business database applications, if not managed appropriately it can result in storage islands that are hard to manage and upgrade scattered throughout the organization.

To address the growth in unstructured data, large enterprise data centers, IPDCs, and service providers are



**WHILE DATA IS
GROWING
RAPIDLY, IT
BUDGETS ARE
RELATIVELY
SHRINKING
AND IT
MANAGERS
ARE ASKED TO
DO MORE
WITH LESS.**

¹ IDC press release: “IDC Predicts 2012 Will Be the Year of Mobile and Cloud Platform Wars as IT Vendors Vie for Leadership While the Industry Redefines Itself”

SCALE-OUT STORAGE

evolving to overcome the limitations of traditional storage infrastructures.

The scale-out storage architecture, shown in Figure 2, implemented using standard high-volume servers, can meet the needs of these evolving data centers. Solutions from many vendors implement the scale-out storage architecture to meet a variety of storage usage models including backup and archive, large object storage, high-performance computing, and business analytics.

As shown in Figure 2, the scale-out storage architecture has three participants:

- **INTELLIGENT STORAGE CLIENT** that understands the two-step access protocol to discover the data it needs
- **METADATA SERVER** that maintains the map of the entire storage landscape
- **STORAGE NODE** that provides the storage media for the data




FIGURE 1. TRADITIONAL ENTERPRISE STORAGE

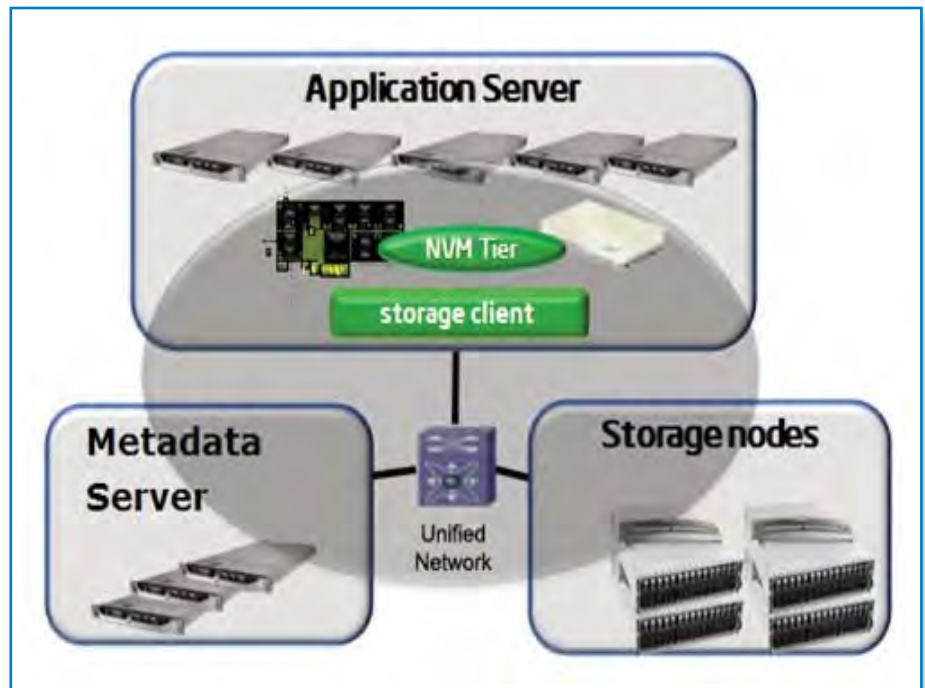


FIGURE 2. SCALE-OUT STORAGE ARCHITECTURE

CONVERGED PLATFORM

All three participants are implemented using appropriate high-volume servers to meet the requirements of specific applications and usage models. The typical interconnect in such an architecture is the unified network that carries both the application and storage data on the same wire, reducing cabling infrastructure and management.

As the industry shifted toward the scale-out storage architecture, the typical high-volume standard servers in the past lacked some of the storage features which are mandatory to meet certain reliability, accessibility, and security (RAS) and SLA requirements imposed by applications.

Intel Corporation's platforms include a number of features that enable previously unavailable storage-specific features on the standard server platforms. The resulting converged platform can serve the standard compute needs of IT organizations while enabling them to deploy large

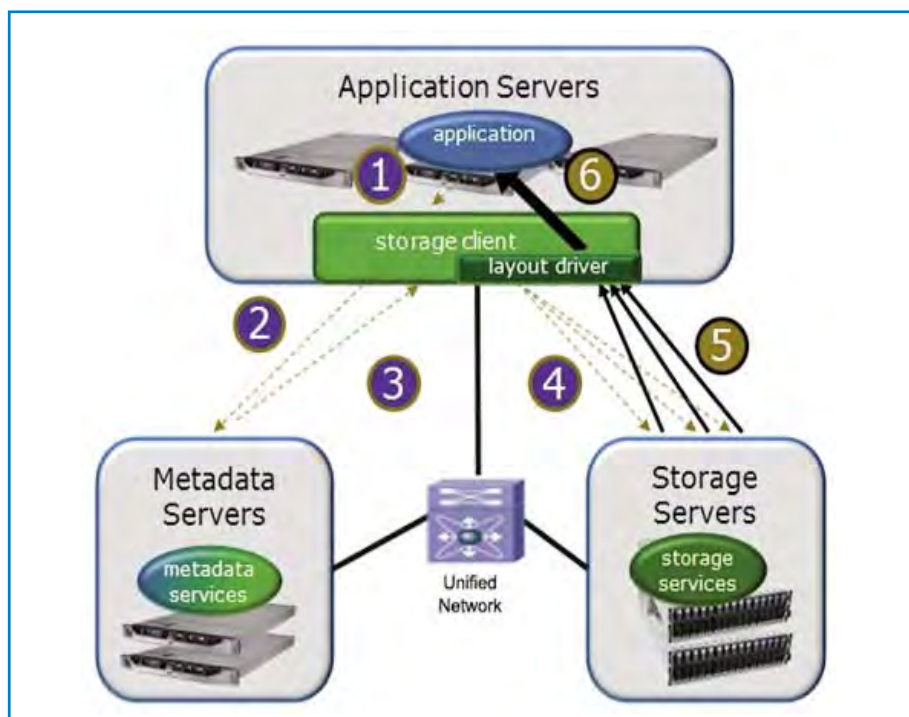


FIGURE 3. SCALE-OUT STORAGE READ OPERATION

amount of unstructured data in a scale-out storage architecture.

Storage-specific features now are enabled in standard server platforms include:

- **ASYNCHRONOUS DRAM REFRESH (ADR).** In a storage application, if software-based RAID was deployed using the host processors, this feature enables battery-powered write-back cache for accelerated RAID perform-

ance plus protection against data loss due to power failure.

- **THE NON-TRANSPARENT BRIDGE (NTB).** Available in the Intel® Xeon® processor E5 family, it enables active-active failover cluster implementation.
- **THE CRYSTAL BEACH DMA ENGINE (CBDMA)** In the Intel Xeon processors E5 family, it allows off-loading of RAID parity calculations, freeing up the host processor for additional compute tasks.

- **NEW INSTRUCTIONS** such as AES-NI can significantly improve certain storage workloads such as inline data encryption and decryption.
- **INTEL® ETHERNET CONTROLLERS** enable a unified network infrastructure that can carry both application and storage data on the same wire.
- **INTEL® SOLID STATE DRIVES** provide storage hierarchy to meet needs of high I/O requirements of many applications.
- **INTEL® INTELLIGENT STORAGE ACCELERATION LIBRARY (INTEL® ISA-L)** accelerates many storage specific algorithms, extracting even more performance out of the storage infrastructure.

The scale-out storage architecture lends itself to a variety of different usage models. Figure 3 shows a typical I/O operation in the scale-out storage architecture to make it easier to understand the roles and functions of the three participants described earlier.

READ OPERATION

In the read operation, we assume that an application is requesting to read an object (data) that was not read previously. This read operation will be carried out in two steps, a discovery phase and a data exchange phase.

In the discovery phase:

1. **AN APPLICATION INITIATES** an object (data) read operation to the storage client.
2. **THE STORAGE CLIENT INQUIRES** about the location map of the stored object to the metadata server.
3. **THE METADATA SERVER SENDS** the location map of the requested object to the storage client. It may reveal multiple storage servers across which the stored object is striped.

In the data exchange phase:

1. **THE STORAGE CLIENT** sends read requests to all the storage servers

identified in the location map to read the stored object.

2. **THE STORAGE SERVERS** respond with partial content of the stored object back to the storage client. All storage servers respond in parallel to the storage client.
3. **THE STORAGE CLIENT** passes the object to the requesting application.

If an object was already read during a past operation, the storage client skips the discovery phase of the operation discussed above and accesses the storage nodes directly to retrieve an object. During the data exchange phase, the storage client and storage server exchange data in parallel, resulting in very high throughput. The scale-

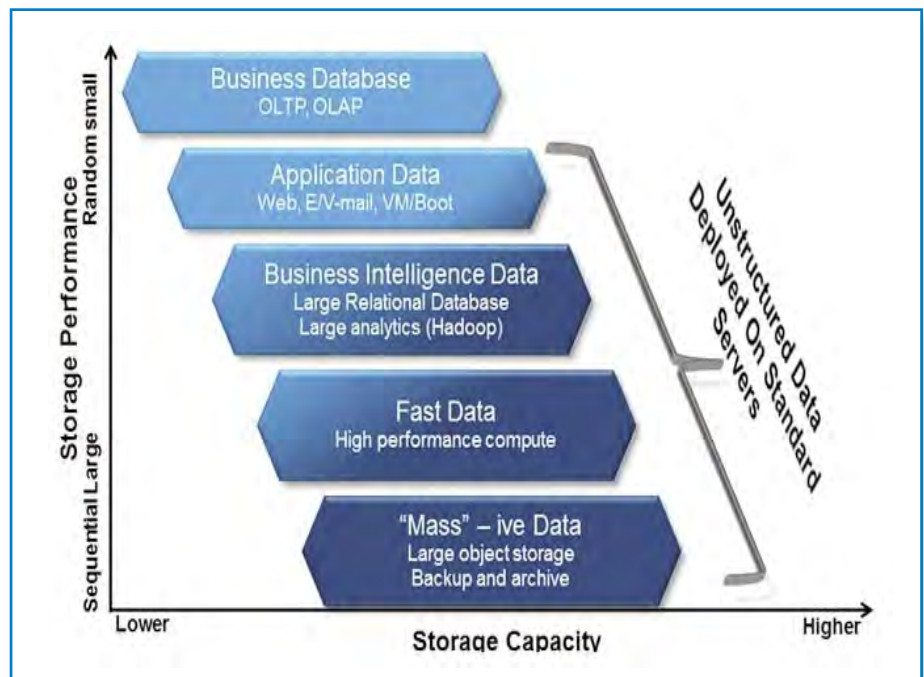


FIGURE 4 STORAGE USAGE MODELS

out storage architecture deployed on standard high-volume servers can serve variety of applications that operate on the unstructured data (Figure 4).

Today's data center has a variety of usage models, which Figure 4 maps by performance versus capacity scale. The higher end of the capacity scale (X-axis) is characterized by higher storage requirements for a specific usage model. The higher end of the performance scale (Y-axis) is characterized by the small random I/Os.

As shown in Figure 4, on the upper-left corner reside the business databases, characterized by small, random I/Os and relatively smaller storage space requirements. On the bottom-right end of the scale resides the massive data, characterized by large, sequential I/Os and petabyte scale capacity. All other usage models, reside somewhere in between on the performance versus capacity scale, as shown in the figure.

A rich ecosystem provides storage software to enable deployment of standard server based solutions for the various storage usage models. For example, the scale-out storage architecture with erasure code is suitable for large object stores, backup, and near-line archives. Solutions based on the Lustre* parallel file system are suitable for high-performance computing environments. Hadoop*-based business analytics applications provide real-time insight into and interpretation of large amounts of machine- and human-generated data.

The scale-out storage architecture is deployed using standard high-volume servers with appropriate storage software. It is suitable for a variety of applications, spanning a vast region on the performance and capacity scale. Since the scale-out storage architecture enables capacity and performance to be scaled independently in small increments gives IT managers better control of their budgets while affording flexi-

bility to meet the performance and capacity demands of their data centers. Thanks to the evolution of the scale-out storage architecture deployed on standard servers, IT managers around the world can manage the explosive growth in unstructured data with much smaller budgets, while meeting wide range of requirements including stringent SLAs, longer retention periods, and higher performance.

As an example, Intel's IT department has achieved capital savings of USD 9.2 million by using a scale-out storage solution in combination with intelligent Intel® Xeon® processors and technologies such as data deduplication, compression, encryption, and thin provisioning. Intel enjoyed these savings while supporting significant capacity and performance improvements. (See the white paper [Solving Intel IT's Data Storage Growth Challenges](#) for details.)

Learn more about cloud storage technologies and tomorrow's cloud [here](#).

[Back to Contents](#)

HADOOP

A Distributed, Parallelized Platform for Handling Large Data

Jeffery Krone

Vice President of Research and Development and Co-Founder
Zettaset, Inc.

jkrone@zettaset.com



Hadoop is a new paradigm for how enterprises store and analyze large data sets, based on the Google File System* and MapReduce* frameworks. Apache Hadoop* is an open source project under the Apache Software Foundation.

FLEXIBLE, SCALABLE STORAGE

Hadoop is essentially a highly scalable enterprise data warehouse that can store and analyze any type of data. It enables distributed, parallelized processing of large data sets across clusters of computers on commodity hardware. Designed for flexibility and scalability, Hadoop has an architecture that can scale to thousands of servers and petabytes of data.

Hadoop is batch-oriented (i.e., high throughput and low latency) and strongly consistent (i.e., data is always available).

Today, there are 1.8 zettabytes of data, 80 percent of which is unstructured. It's expected that by 2017 there will be 9 zettabytes of data, of which 7.2 zettabytes will be unstructured or semi-structured data. Existing solutions such as traditional legacy databases (e.g., MSSQL*, Oracle*, DB2*) and data warehouse products are very good at handling online analytical processing (OLAP) and online transac-

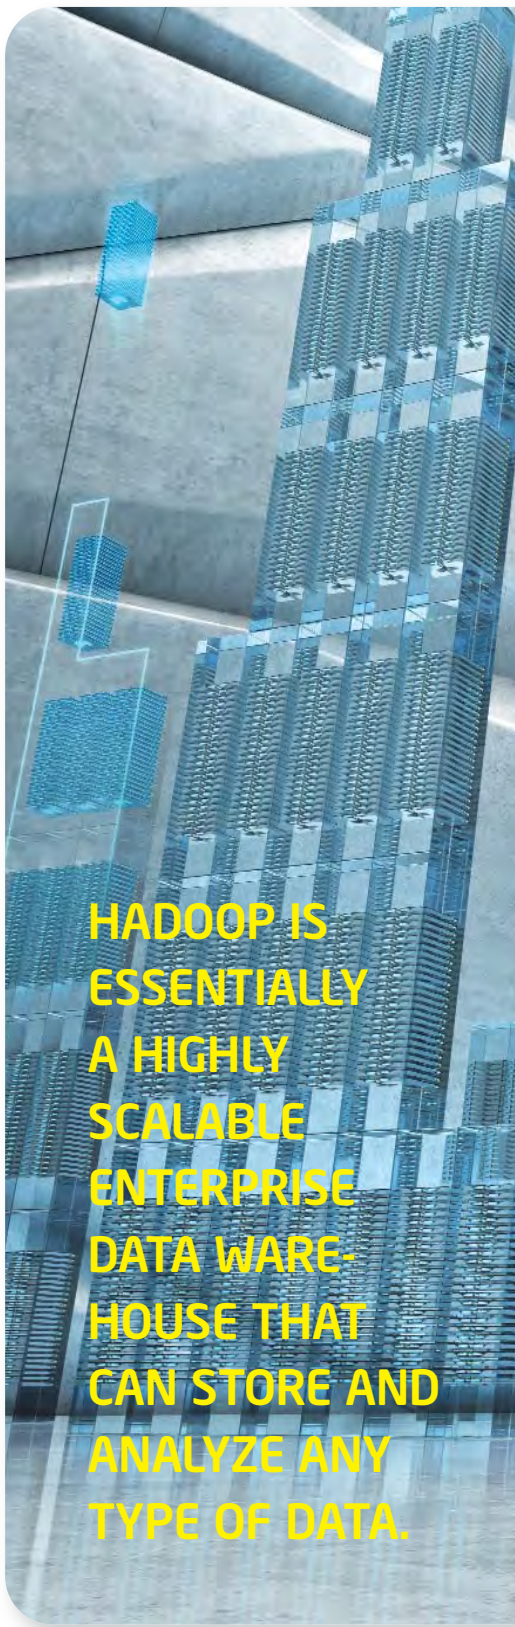
tion processing (OLTP) workloads over structured data.

Hadoop was designed to solve a different problem, the fast and reliable analysis of massive amounts of both structured and unstructured data such as video, audio, social media, and free-form notes.

Many companies are already deploying Hadoop alongside their legacy systems to aggregate and process structured data with new unstructured or semi-structured data to derive actionable intelligence.

Hadoop has two major subsystems:

- **HADOOP DISTRIBUTED FILE SYSTEM** (HDFS) allows reliable storage of petabytes of data across clusters of machines.
- **MAPREDUCE** is a software framework for writing applications that process very large data sets in parallel across clusters of machines. It enables the user to run analytics across large blocks of data.



HADOOP IS ESSENTIALLY A HIGHLY SCALABLE ENTERPRISE DATA WAREHOUSE THAT CAN STORE AND ANALYZE ANY TYPE OF DATA.

MEETING THE CHALLENGES

Essentially, MapReduce is a programming paradigm that divides a task into small portions and then distributes it to a large number of nodes for processing (map). The results are then summarized in the final answer (reduce).

Other components of the Hadoop ecosystem include:

- **HBASE***, based on Google's BigTable*, a non-relational, scalable, fault-tolerant database that sits on HDFS.
- **ZOOKEEPER***, a centralized service that maintains configuration information (i.e., roles of various servers on the cluster), naming, and group services such as leadership election (i.e., electing a new backup server among a group of servers in the cluster).
- **OOZIE***, a job scheduler that works in conjunction with Hadoop. Essentially, Oozie is a workflow engine that enables users to schedule a series of jobs within a Hadoop framework.
- **HIVE*** abstracts the complexities of writing MapReduce programs.

Hive enables a user to write SQL*-like queries, which are then converted into MapReduce programs.

This allows a user to manipulate data on a Hadoop cluster (e.g., select, join, etc.). Most people who are familiar with traditional relational databases will find it easy to program in Hive, which is very similar to standard ANSI SQL*.

- **PIG*** is a high-level procedural language for querying large data sets over Hadoop. Pig abstracts the complexities of writing MapReduce programs.

CHALLENGES IN DEPLOYING AND BUILDING A HADOOP CLUSTER

As you begin to build your Hadoop cluster, you will probably run into several challenges including:

- Provisioning and monitoring
- High availability
- Backup
- Scheduling jobs
- Complex writing of MapReduce programs
- No user interface

- Weak security
- Limited import and export capabilities

PROVISIONING AND MONITORING

You must provision/install all of your storage, server, network, and Hadoop services. This will entail distributed logging, configuration management, an alerting mechanism, failure recovery, service migration, performance monitoring in the cluster, and automatic installation and provisioning of servers to scale to potentially thousands of servers within a cluster.

HIGH AVAILABILITY

Supporting large-scale clusters requires fault-tolerant software. While Hadoop is very resilient to the failure of individual nodes, the primary name node is a single point of failure. If the primary name node fails, then no file system operations can be performed. If the primary name node is unrecoverable, you could lose all file system metadata. This would make it impossible to reconstruct the files stored on the cluster.

FILLING THE GAPS

To protect against primary name node failure, and to keep other key services within Hadoop from failing, you will need to have hot standby servers for all key services of Hadoop (e.g., primary name node, job tracker, task tracker, Oozie, Hive Metastore, data node, Kerberos*, secondary name node, and Zookeeper). The support of your Hadoop cluster, which could range from a few machines to potentially thousands of nodes, requires at least one or more dedicated administrators.

If the cluster is not self-healing, and is not designed appropriately for redundancy, automatic failover, or proactive monitoring, you would need a substantial support staff on call 24x7, which could cause your OpEx costs to spiral out of control.

BACKUP

There are currently no tools to help you backup a Hadoop cluster of up to thousands of terabytes of data. Performing incremental or full backups at this scale is beyond the capability of existing backup tools.

SCHEDULING JOBS

Currently, there is no system available for scheduling periodic MapReduce

jobs or Hive or Pig workflows on Hadoop. You can only submit jobs for immediate execution.

COMPLEX WRITING OF MAPREDUCE PROGRAMS

MapReduce programs are difficult to write. You will need to integrate Hive and Pig to abstract the complexities of MapReduce.

NO USER INTERFACE

Hadoop does not currently have a user interface. You must use the command line to issue requests to the system.

Therefore, it can be hard to monitor and manage the Hadoop ecosystem.

WEAK SECURITY

Hadoop security is fairly weak. Hadoop has partially adopted Kerberos authentication, but many services remain unprotected and use trivial authentication mechanisms.

LIMITED IMPORT AND EXPORT CAPABILITIES

Hadoop has some minimal support for reading and writing basic flat files. Users are required to roll out their own import and export formats.

GAPS IN HADOOP AND HOW TO ADDRESS THEM

PROVISIONING AND MONITORING

Hadoop is a very complex ecosystem that is hard to install, provision, and monitor. The Zettaset* (ZTS*) data platform will alleviate this problem through its ability to automatically install, configure, and provision servers within a Hadoop cluster. Functionality will be available in the next few months.

Essentially, remote agents will install a selected Hadoop distribution (e.g., CDH*, IBM Big Insights*, Apache), and ZTS packages on the nodes within the Hadoop cluster. A centralized configuration depository (Zookeeper) will download a Hadoop role (e.g., primary name node, secondary name node, or task tracker) to a specific node on the cluster. Once the configuration files for a specific role are downloaded to a node on the cluster, the appropriate Hadoop services will be instantiated and the node will assume that role.

Through the ZTS user interface, you can change the role of a particular node as required. The ZTS platform also has the capability to monitor, start, and stop key Hadoop services through the user interface.

HIGH AVAILABILITY

The current Hadoop distributions don't have a failover mechanism for all critical Hadoop services. Some of the Hadoop distributions, such as Cloudera*, do address manual primary name node failover. However, none of the other key Hadoop services are backed up.

In contrast, ZTS has a straight-forward way to handle high availability that you can apply to all key Hadoop services (e.g., primary name node, secondary name node, job tracker, task tracker, Oozie, Hive Metastore, network time protocol, and Kerberos) and other miscellaneous services (MongoDB*, SQL Server*). The backup mechanism is 100 percent automated and supports both stateful and stateless failover.

A stateful failover (Figure 1) is when a key service such as the primary name node fails. Since data is written

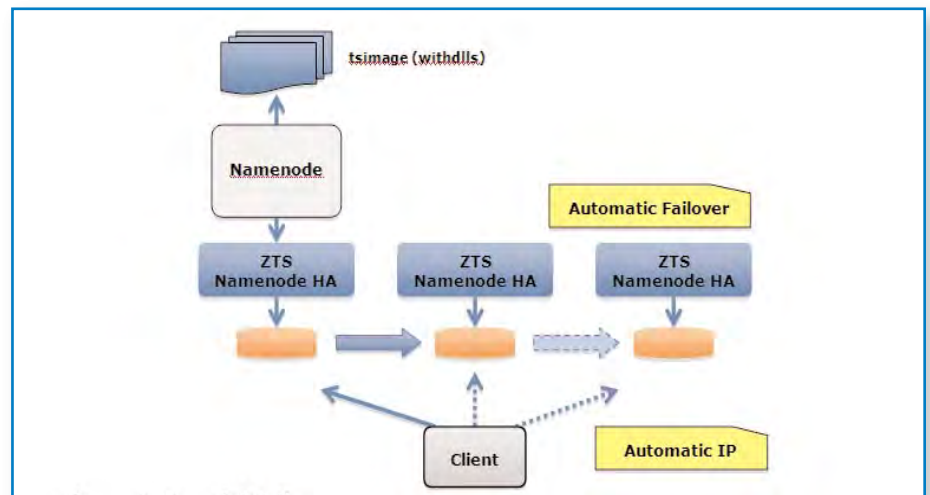


FIGURE 1. STATEFUL FAILOVER

to both the primary and backup name nodes concurrently, there is no data loss. The new primary name node is fully synchronized from a data perspective with the old primary name node. The new primary name node assumes the IP address of the failed primary name node and is brought up automatically. Therefore, domain name system (DNS) updates are not required upon failure.

A stateless failover is when a stateless service such as task tracker fails and is automatically restarted. No state information is maintained.

BACKUP OF THE CLUSTER

Since a Hadoop cluster can range from a few terabytes to tens of petabytes of data, it's hard to backup an entire cluster or even a portion of the cluster. A few

Hadoop vendors such as MapR address this issue by allowing users to take snapshots of their clusters and perform incremental backups as required. Zettaset is adding the ability to take snapshots of the cluster in the near future. Currently, Zettaset provides the ability to back up a cluster to another cluster in a different data center. However, state information between the clusters is not maintained (ie, data is not replicated between the clusters continuously).

SCHEDULING WORKFLOWS

The Hadoop distributions, as they stand today, do not have an intuitive graphical user interface for defining and scheduling Hadoop workflows. The only mechanism currently available for defining workflows is Oozie. However, Oozie is based on .xml* files that have to be manually set up by users. All files associated with various jobs (MapReduce, Hive, and Pig) must

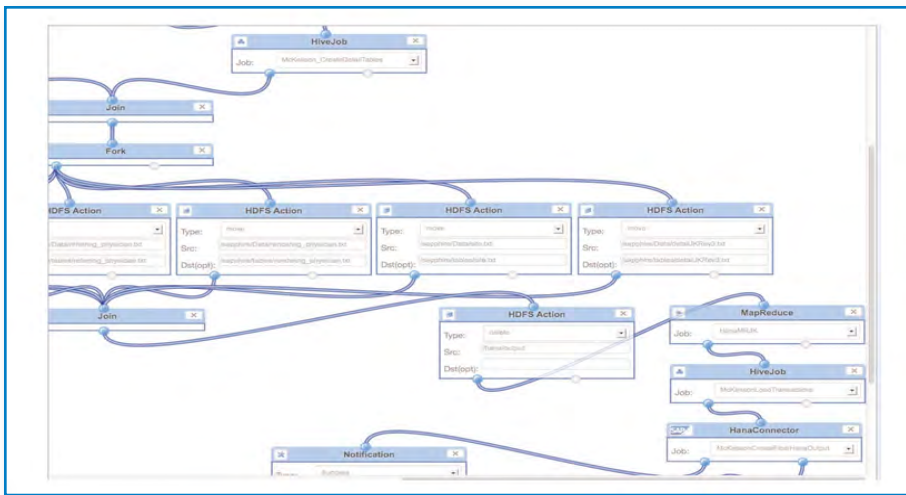


FIGURE 2. TABLET ID AS A BANDED PLATFORM OF THE PRIVATE CLOUD



FIGURE 3. ZETTASET USER INTERFACE




FIGURE 4. ZETTASET USER INTERFACE MANAGES AND MONITORS HADOOP COMPONENTS

grams, and the scheduling of workflows (Figure 2).

INTEGRATING HIVE AND PIG INTO THE ZTS PLATFORM

To abstract the complexities of writing MapReduce programs, the Hadoop community has created Hive and Pig. Zettaset has integrated Hive into its user interface (UI) (Figure 3). This enables users to create and configure Hive jobs using the Zettaset UI. As an added feature, Zettaset will verify that the syntax of the Hive job is valid and display any errors that in the Hive job syntax. Zettaset has also embedded the Hive console within its UI. Zettaset has plans to include Pig within its UI in the near future.

USER INTERFACE FOR HADOOP

Currently, the various Hadoop distributions do not have an associated user interface, making it hard to manage Hadoop components and monitor key Hadoop services. To simplify management of the Hadoop cluster, Zettaset has created a user interface that sits on top of Hadoop (Figure 4). The Zettaset user Interface has these capabilities:

be manually copied to the appropriate directory structures. Zettaset is the first company to create a simple, intuitive, drag-and-drop interface for Oozie. The ZTS plat-

form supports most of the Oozie functionality such as flow control modes (e.g., fork and join), HDFS file actions (e.g., move, delete, and create), system notifications (e.g., email notifications of job status), the execution of hive and MapReduce pro-

- **CAN MONITOR**, start, and stop critical Hadoop services (e.g., primary name node, job tracker) on each node of the cluster.
- **THE ZETTASET UI** has an HDFS browser, which enables a user to manage the HDFS file system (e.g., import and export files from network-attached file systems; view, create, or delete directories and files within HDFS; and assign permissions to HDFS directories and files).
- **USERS CAN CREATE**, schedule, and configure Hive jobs through an intuitive user interface.
- **USERS CAN CREATE**, schedule, map, configure, and reduce programs through an intuitive user interface.
- **USERS CAN** easily create complex Hadoop workflows.
- **ADMINISTRATORS CAN** assign role-based security to users and groups (HDFS permissions, MapReduce, Hive, and workflows) through the user interface.
- **THE ZETTASET USER INTERFACE** will, in the near future, allow users to manage multiple Hadoop clusters from a single UI.

HADOOP SECURITY

The current Hadoop distributions have very limited security mechanisms (ie, limited Kerberos authentication and trivial authentication). For instance, if you execute a Hive workflow, MapReduce program, or Oozie, a root user becomes the proxy to execute the job. Therefore, any user who has access to the system can execute any of those tasks.

To enhance the security on Hadoop, Zettaset has implemented role-based security, which applies to HDFS permissions, MapReduce, Hive, and Oozie workflows. This ends the ability of users or groups to access data or execute jobs and workflows for which they don't have access.

Role-based security is now implemented in the Zettaset user interface. In the near future, it will also be available in the Zettaset API and command-line interface. Also, Zettaset has removed the complexities associated with integrating Kerberos authentication with the various Hadoop components by automatically implementing Kerberos authentication across the entire Hadoop ecosystem.

In the near future, Zettaset will:

- Use hardware encryption to encrypt data residing in HDFS and communications between various Hadoop services
- Implement a comprehensive audit log that will capture all activities performed by various users and groups
- Be able to synchronize Zettaset role-based security with Microsoft Active Directory*, Lightweight Directory Access Protocol* (LDAP*), and UNIX* file security

Figure 5 shows the Hadoop security set-up.

SUPPORT FOR IMPORTING AND EXPORTING DATA INTO HADOOP

Hadoop support for importing and exporting data is limited. Users must write and roll out their own import and export formats. Zettaset will embed Flume* and Sqoop* as its extract, transform, and load (ETL) mechanisms. This allows the ZTS platform to connect to any relational database via a Sqoop (e.g., DB2*, Oracle, MSSQL*, PostgreSQL*, MySQL*) and import any type of log file (e.g., Web logs, firewall logs, error log files, and SNMP) from any machine via a distributed agent infrastructure [Flume].




FIGURE 5. HADOOP SECURITY

RUNNING HADOOP IN THE CLOUD (VIRTUAL SERVER INFRASTRUCTURE)

There are several advantages to running Hadoop in the cloud:

- **A HADOOP CLUSTER** can be installed very quickly within a cloud environment.
- **THE HADOOP CLUSTER** is centralized and can be easily accessed via the Internet.
- **A USER CAN EXPAND OR CONTRACT** the Hadoop infrastructure easily by instantiating or de-instantiating node instances as required.
- **IF HADOOP NODE IMAGES FAIL**, it is very easy to recreate them in a virtualized cloud environment.
- **THROUGH ELASTIC IPS**, the implementation of high availability for the Hadoop cluster is plausible.

There are also disadvantages to running Hadoop in a cloud computing environment:

- **IF YOU MNEED YOUR CLUSTER** to be operational 24x7, Amazon's per-CPU costs become very expensive, easily comparable to running your own server in a data center when you factor in power and network bandwidth.
- **HADOOP IS** RAM- and storage-intensive. You will need a cluster of large, or even extra large, high-CPU instances for Hadoop to perform well.
- **AMAZON STORAGE** is problematic. The S3* file system is slow and has no locality, but it is durable. Therefore, network latency will become a factor.
- **THE EBS* FILE SYSTEM** is faster than S3 but not as reliable, although compatibility is still comparable to a real-world SAN. However, EBS is also more expensive if you perform large-scale analytics, since it charges for both storage and for each I/O operation. Most analytic jobs read and write entire data sets.
- **VIRTUALIZED, SHARED** network cloud environments such as Amazon EC2* often experience periods of very high latency during peak traffic conditions, even for internal communications within a cluster. This can cause severe problems for HBase and Zookeeper, which rely on timely network responses to determine that a machine is online and operational.

- **SINCE YOU DON'T MANAGE** Amazon's network, you can't perform traditional Hadoop optimizations for maximizing data locality and performance.
- Amazon charges for all external I/O (i.e., packets into and out of your geographic region). Factor this cost in; if you're importing terabytes per month, it can be prohibitive.

Overall, Hadoop is an excellent choice for enterprises that need to process massive data sets that include both structured and unstructured information. But current Hadoop distributions are not enterprise-ready. They can be hard to install, provision, and manage with multiple single points of failure.

Using the ZTS platform, which runs on top of any Hadoop distribution, can eliminate many of the gaps in Hadoop. The ZTS platform makes Hadoop enterprise-ready and makes it easy to install, provision, manage, and monitor large Hadoop clusters on any type of hardware or in a cloud environment.

To learn more about Hadoop, visit www.hadoop.apache.org.

To learn more about the ZTS platform, visit www.zettaset.com.

Back to Contents

NEXT-GENERATION CLOUD MANAGEMENT

The CompatibleOne* Project

Jean-Pierre Laisné,

CompatibleOne Coordinator and Open Source Strategy
and OW2 Chairman, Bull

jean-pierre.laisne@bull.net

Iain James Marshall

CompatibleOne CTO and Technology Strategy Manager, Prologue

ijmarshall@prologue.fr

Parviz Peiravi

Editor in Chief, *Journey to Cloud*,

Intel Corporation

parvis.peiravi@intel.com

In the nearly four years since its introduction, cloud computing has moved from hype to key IT strategy. Enterprises of all sizes have turned to cloud computing for savings, scalability, and flexibility—despite such concerns as security, privacy, data management, performance, and governance.

MATCHING NEEDS TO SERVICES

Cloud is not only driving significant changes IT business models, it is also driving structural changes in IT, forcing a more holistic approach and a new operating model that enables IT to source and recommend services the enterprise values, in real time and with guaranteed performance and reliability.

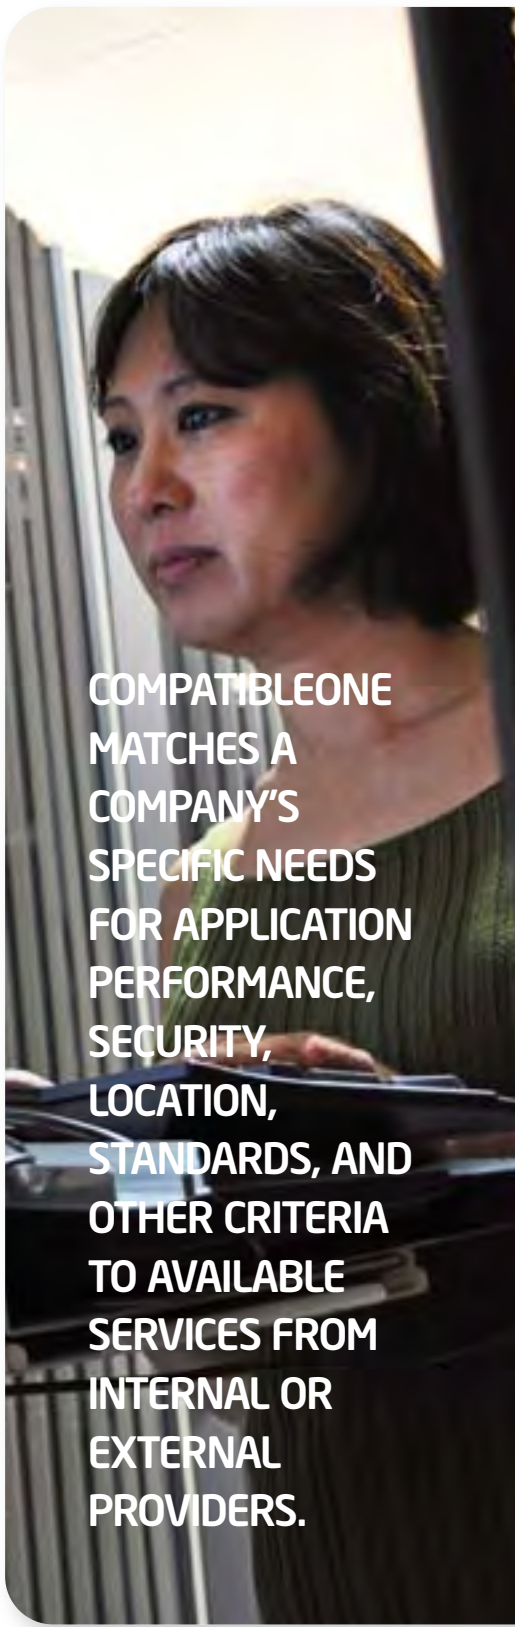
IT organizations are positioning themselves as central service brokers for the enterprise. They are creating organizational structural changes in unprecedented response to environmental changes, restructuring the organization around a new model that will enable faster adoption of services through a governed and balanced sourcing strategy, enabling business units to differentiate their value and gain a competitive edge.

The goal of IT in this new environment is to share resources among the cloud service consumers and cloud service providers in the cloud value chain in a model called IT as a service. The service broker in this model provides an integrated ecosystem that streamlines the

sourcing and provisioning of services from multiple providers and exposes them to enterprise users. The service broker requires a complex infrastructure with elasticity, automation, enhanced monitoring, and interoperability with both the existing system and new services. Those issues have been barriers to the adoption of IT as a service in the enterprise environment.

The CompatibleOne* project is an effort to address the complex barriers to cloud adoption and enable organizations to realize the full potential of cloud services. CompatibleOne matches a company's specific needs for application performance, security, location, standards, and other criteria to available services from internal or external providers. It then creates and deploys the appropriate service within the cloud, ensuring that cloud services work together to meet the company's needs.

CompatibleOne is an open source middleware (or cloudware) that gives users of cloud services interoperability with one or more cloud



**COMPATIBLEONE
MATCHES A
COMPANY'S
SPECIFIC NEEDS
FOR APPLICATION
PERFORMANCE,
SECURITY,
LOCATION,
STANDARDS, AND
OTHER CRITERIA
TO AVAILABLE
SERVICES FROM
INTERNAL OR
EXTERNAL
PROVIDERS.**

DELIVERING BUSINESS VALUE

service providers. Since it's open source, this cloudware may be used by other initiatives. It's complementary to many other open source projects such as OpenStack* and OpenNebula*, with which CompatibleOne shares the aim of helping to foster a sustainable ecosystem based on the open cloud concept, where cloud infrastructures are based on free software and interfaces and open standards and data formats. CompatibleOne is used both by partners in the project and by a wider community, including French and European projects.

The main characteristics of CompatibleOne are:

- **INTEROPERABILITY.** A way to integrate and aggregate services provided by all types of private, community, or public clouds.
- **PORTABILITY.** A way to move workloads from one cloud provider to another.
- **REVERSIBILITY.** A way for the user to recuperate data and processes.

- **RESPECT FOR SECURITY AND QUALITY OF SERVICE ON ALL-LAYERS OF THE CLOUD.**

Infrastructure, platform, applications, and access.

To develop applications with all these characteristics, CompatibleOne created a DevOps blueprint of cloud computing and all its resources, including APIs and necessary protocols. This work has highlighted the importance of modeling, both to foster interoperability and to provide an abstraction of services regardless of their providers. It can also facilitate cooperation among application developers, architects, and operators.

Based on a service architecture, CompatibleOne offers a new way to provision and distribute workloads in the cloud, starting with the needs of users (e.g., end users, IT department operators, system integrators, and developers). Because it functions as an intermediary between consumers of services and complex offerings, because it allows the integration of a variety of services, and because it facilitates their

COMPATIBLEONE PARTICIPANTS

Launched in 2010,

CompatibleOne is co-funded by Fonds Unique

Interministériel, Région Ile de

France, Conseil Général des

Yvelines, and Mairie de Paris

and supported by OSEO,

Systematic, and Pôle SCS.

Participants in the project

include ActiveEon, Bull,

CityPassenger, eNovance,

INRIA Rhône-Alpes, INRIA

Méditerranée, Institut

Télécom, Mandriva,

Nexedi, Nuxeo, OW2,

Prologue, and XWiki.

For more information, visit

www.compatibleone.net.

selection, CompatibleOne is essentially a cloud service broker, as defined by the NIST reference architecture and research by Gartner and Forrester.

FLEXIBLE AND ROBUST

USAGE MODELS

The flexibility and robustness of the CORDS model and the ACCORDS platform allow users to fully independently operate all resources provided by a heterogeneous mix of providers (e.g., OpenStack, Azure, OpenNebula, and Amazon), which prevents the problem of vendor lock-in.

In the same way, users have transparent access to a heterogeneous mix of services provided by infrastructure as a service (IaaS) or platform as a service (PaaS). For example, CompatibleOne makes it possible to port images to any hypervisor and, at provisioning time, it will produce an image compatible with the hypervisor used by the selected service provider.

The CompatibleOne model makes this management of heterogeneity possible. CORDS allows for modeling of any cloud computing service to enable it to be provisioned by any provider. It offers a complete abstraction of the infrastructure, platform, and service, regardless of the provider. CORDS makes it possible to

conceive and create interoperability.

The CompatibleOne platform allows users access to interoperable clouds now, without waiting for standards to mature.

The CompatibleOne solution isn't unique to a single problem. In fact, it was designed with several cloud computing market segments in mind and aims to provide a comprehensive solution to clearly identified problems in the fields of:

- **TELECOMMUNICATIONS AND INTERNET SERVICE PROVIDER OPERATIONS.** Managing multiple, heterogeneous resource centers for a customer base that includes both corporate enterprise and governmental bodies and also domestic and roaming users.
- **CORPORATE ENTERPRISE INFORMATION TECHNOLOGY MANAGEMENT.** The focus is on optimizing operational costs, flexibility, and adaptability to new marketplace trends, plus secure access to vital corporate resources.
- **BUSINESS APPLICATION VENDOR SYSTEMS.** This requires a comprehensive approach to managing multiple customers, vendors, and service providers. The focus is on cost-effective provisioning of resources with control of cost margins and customer returns.

There are several use cases for CompatibleOne. For example, imagine a private cloud and an IT department at a major company. With the emergence of the cloud, the IT department is clearly in danger of losing control of its IT systems, since internal clients prefer to deal directly with cloud service providers instead of with IT.

CompatibleOne gives this IT department a way to offer, on-site, a private cloud management service that can meet the demand from internal clients while retaining control of architecture and responsibility for negotiations with providers.

This means the IT department will be able to offer a catalog of servic-

PRIVATE CLOUD MANAGEMENT

es, associated with a list of certified-providers, depending on selection criteria, and in full compliance with the company's security policy. The selection criteria can be performance requirements, quality of service requirements, or location criteria to meet the sovereignty needs of the enterprise (e.g., "I would like my data to be stored only in Europe and for only my users to have access to this part of the cloud"). This enables the IT department to satisfy its users and negotiate optimal contract terms and service level agreements (SLAs) with various providers, in compliance with the business and legal environment of the company. It can also mix and match computing and storage services from different providers, depending on economic criteria. With CompatibleOne, the IT department can concentrate on its core business and improve its range of services while controlling costs and maintaining high performance standards. By customizing the CompatibleOne platform, the IT department can offer its users (or business units) value-added services according to the strategic orientation of the enterprise. For example, it could plot service usage on a calendar and

plan for supplementary resources during activity peaks.

In another example, IT could develop a community cloud that combines public and private cloud services. Shared by users with common cultural, commercial, or organizational areas of interests, this cloud could offer access to the best service providers according to criteria such as professional specialization or geographical location. Using the CompatibleOne platform to manage intermediation and integration of heterogeneous public and private services gives this community of users access to the best services at the best prices, combining services to satisfy shared needs while also ensuring secure access to resources.

The final use case demonstrates the flexibility of the CompatibleOne model and platform: high-performance computing (HPC). Consider the example of a scientific computing center that manages tens of thousands of servers and needs to offer its users adapted computing capacities. The managers of the computing center need to:

- **MEET THE NEEDS** for massively parallel computing that requires a dedicated infrastructure
- **PROVIDE MORE ECONOMICAL** (i.e., elastic) resources such as virtual clusters managed with OpenStack or OpenNebula, based on standard CPUs or GPUs

With CompatibleOne, they can model these various types of infrastructures, automatically distribute workloads depending on their resource needs, and provision the resources depending on selection criteria (e.g., smart allocation of processing software depending on CPU or shared memory needs). CompatibleOne can completely automate processing distribution over several types of heterogeneous clouds. By extension, because they can aggregate heterogeneous services, the architects of these systems can also design ways to interconnect them securely with public clouds, such as specialized image libraries, or with other private clouds that offer complementary

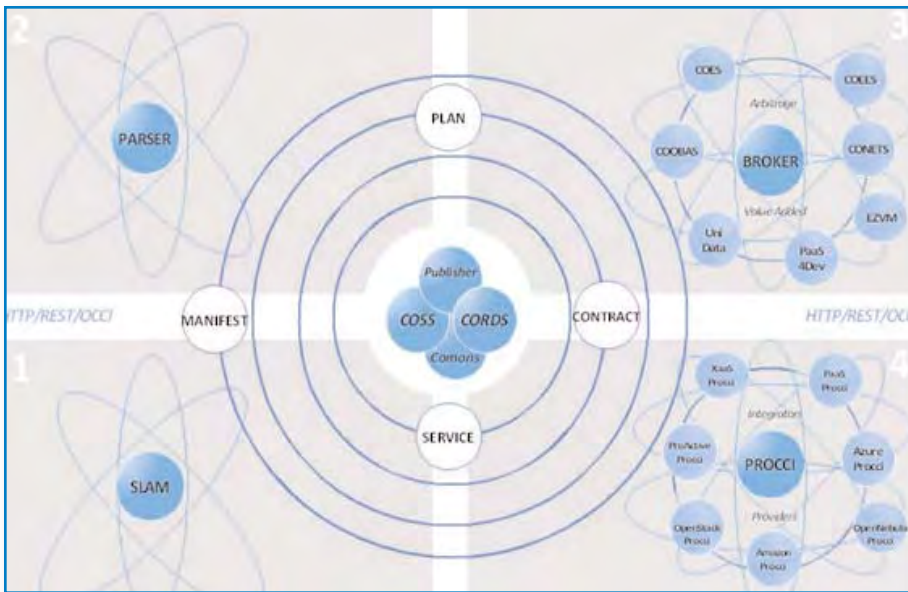


FIGURE 1. ACCORDS PLATFORM

geographical information needed for the calculations.

With its REST architecture, its CORDS model, and its ACCORDS platform, CompatibleOne provides basic services—such as security, monitoring, billing, and brokering—to support use cases for various types of cloud (private, public, and hybrid) and to create a platform adapted to the needs of CIOs, system administrators, operators, and brokers. (The term “brokerage services” may refer to technical services or to an intermediation or broker-type business model, as defined by Gartner and Forrester.)

CompatibleOne also makes it easy to create integrated and innovative value-added services. In short, CompatibleOne

offers integrators and developers a platform that can adapt to their projects.

SOLUTION OVERVIEW

Using the ACCORDS platform has four steps (Figure 1):

- **MANIFEST SUBMISSION.** Requirements for the provisioning of cloud resources are described using the CompatibleOne Request Description Schema (CORDS) and submitted to the system in the form of an XML* document called the manifest.
- **RESOURCE PROVISIONING PLAN:** The ACCORDS Parser validates and processes the manifest, producing a fully qualified resource provisioning plan. This provisioning plan describes in precise detail the

operations to be performed for constructing and delivering the cloud application configuration.

- **PROVISIONING OPERATION.** The provisioning plan can be used at any time to provision the cloud resource configuration as described by the manifest. This provisioning operation is performed by the ACCORDS Broker, working in cooperation with the placement components (COES) and provisioning components (PROCCI) of the platform. Placement means selecting not only the most appropriate provisioning platform type, but also the right commercial collaborator to provision resources. Looking to meet the needs of the use cases, the powerful and flexible algorithms of the placement engine allow their decisions to be based on technical, financial, commercial, geographical, performance, and quality of services considerations.
- **DEPLOYING APPLICATIONS AND HARDWARE.** Finally, heterogeneous provider platforms deploy the applications and hardware required to satisfy the configuration as described by the manifest. When working with predetermined

quotas negotiated in advance, failure of any particular provider or collaborator to deliver is fed back to the placement engine for selection of alternative providers. This allows for not only fail-over management, but also for real-time assessment of quality, both operational and commercial, of all involved parties.

These major operational components are loosely and flexibly interconnected through a collection of service components. This makes it easy to support different usage scenarios by integrating or replacing operation-specific components.

You can implement each operational concept—and component—of the platform as an individual, standalone service management platform for unlimited scalability.

The generic provisioning interface, provided by the components that make up the ACCORDS PROCCI, lets you extend the platform by adding different provisioning components in real time as input manifests. You can use this process to create subse-

quent instances of the platform itself to flexibly meet growing provisioning needs.

The CORDS model (Figure 2) provides a complete and comprehensive set of object-oriented tools and constructions for describing and provisioning cloud resources.

Manifest descriptions of provisioning configurations can include both simple node definitions (in terms of their virtual hardware and application software needs) and more complex provisioning systems (represented by other manifest or class descriptions).

A collection of configuration actions—including instructions for interconnect-

ing nodes, activating monitoring, or invoking usage-specific service methods—describe instances of provisioning. Manifest descriptions can expose interface methods, allowing provisioned instances to cater to particular needs when contributing service components to other client service instances. It's possible to describe nodes so that they contribute their characteristics and services to either single or multiple instances, either respecting or indifferent to the defining manifest or class. In this way, it's possible to define single service components that cater to the needs of either a single manifest or more global communities of provisioning users and customers.

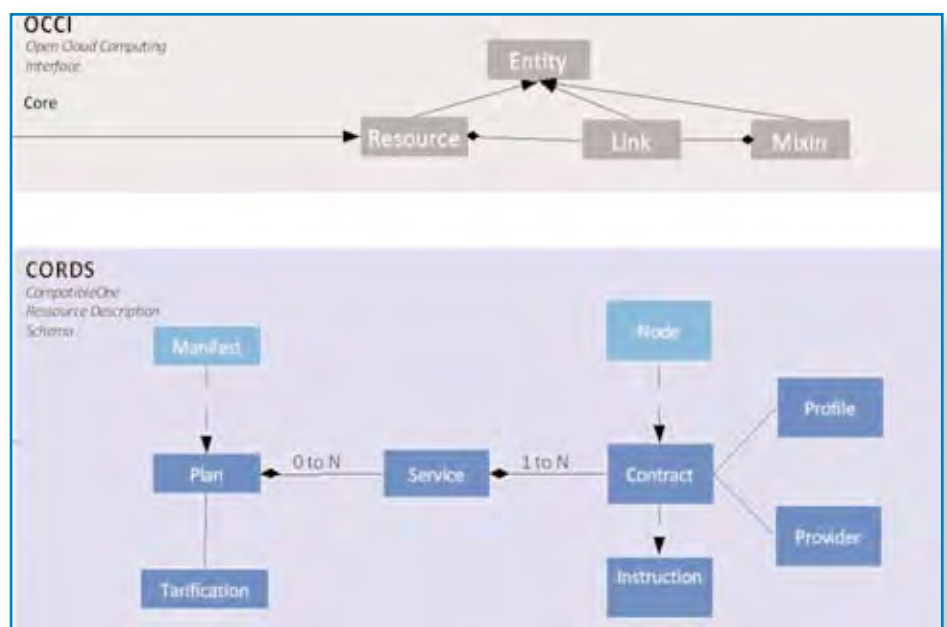


FIGURE 2. CORDS LOGICAL VIEW

These constructions make it possible to use CORDS to describe provisioning in the realm of IaaS, describing precise requirements using simple nodes to deploy application images of all types and uses.

You can use the techniques of complex node description, in conjunction with manifest service instance declaration, to meet the needs for describing and deploying PaaS offerings.

The platform services enable developers to use the platform of their choice (e.g., JavaEE*) to run their applications. Developers aren't forced to port their applications according to specific and proprietary features of a PaaS provider such as AWS Beanstalk* or Google App Engine*.

Once the developer chooses a platform and has it described by a CORDS manifest, CompatibleOne provisions the platform services.

CompatibleOne's approach to this is simple. Once the platform is described by a CORDS manifest, the developer can deploy it on the ACCORDS platform.

A component of the platform, known

as the PaaS Procci, then makes itself known through ACCORDS Publisher as a provisioner of service for deploying a particular type of node. The PaaS Procci also submits two types of manifests for parsing by the ACCORDS Parser:

1. One that describes services the platform offers the end user
2. One that describes the resources the PaaS requires to deliver these services

In this way, the developer can use nodes representing a PaaS service in application manifests, allowing seamless integration with more traditional IaaS nodes.

Extending the PaaS technique, and using interfaces exposed by their provisioned components, the ACCORDS platform can provide symmetrical solutions for dynamically integrating heterogeneous components and services required to construct more complex service systems such as XaaS and BPaaS.

VIRTUAL INSTANCE

The manifest and resulting provisioning plan represent a particular class of

provisioning configuration. The act of provisioning performed by the ACCORDS Broker produces a provisioning control structure known as the Service Graph. This includes the contracts negotiated by the placement engine to satisfy the needs for provisioning. Accompanying each contract is a list of instructions that ensure the configuration and monitoring of the component contract within the particular instance of service (Figure 3).

COMMUNICATION ARCHITECTURE

The CompatibleOne platform communication architecture rests solely on the Hypertext Transfer Protocol (HTTP) operating in a RESTful way, providing a loosely-coupled operational framework. All server components make up the architecture that exposes a standard OCCI interface, describing the collection of service categories offered by each particular component. Each public service category is published through the ACCORDS Publisher component, which assumes the central role of the platform, allowing for service discovery in a dynamic, flexible and extensible way.

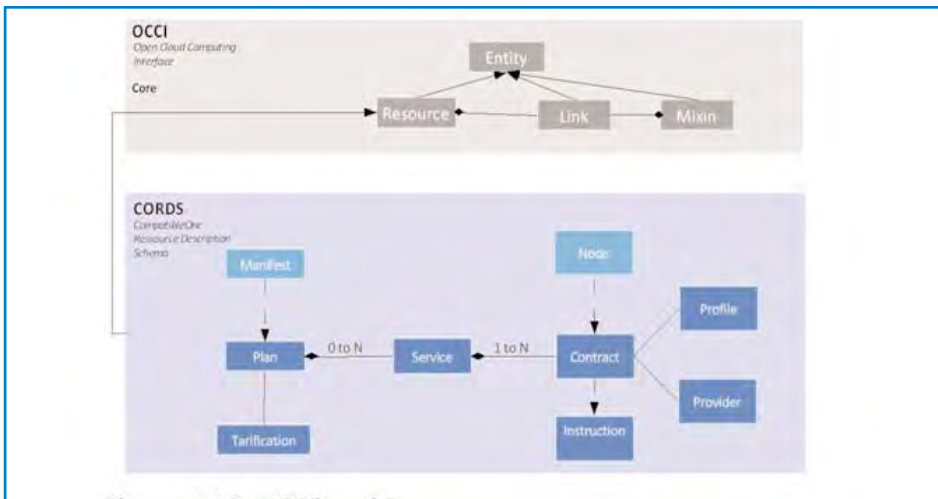


FIGURE 3. CORDS VIRTUAL INSTANCE

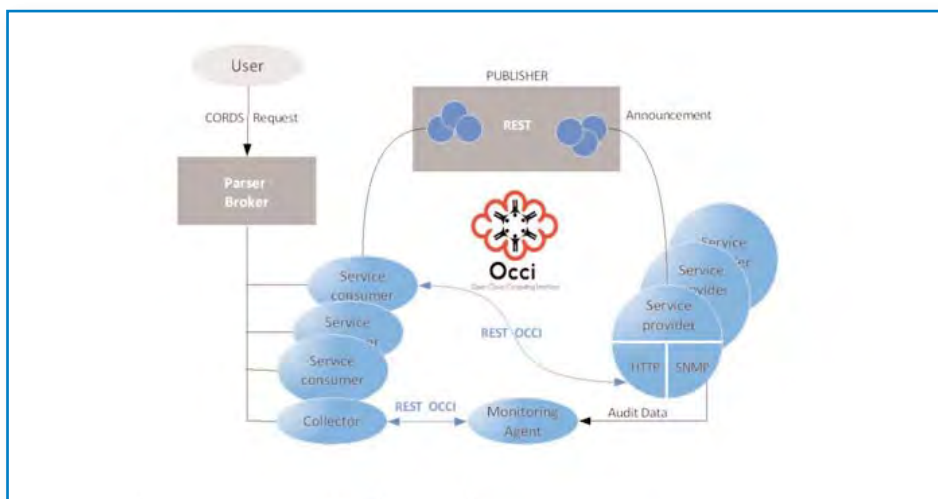


FIGURE 4. COMMUNICATION ARCHITECTURE

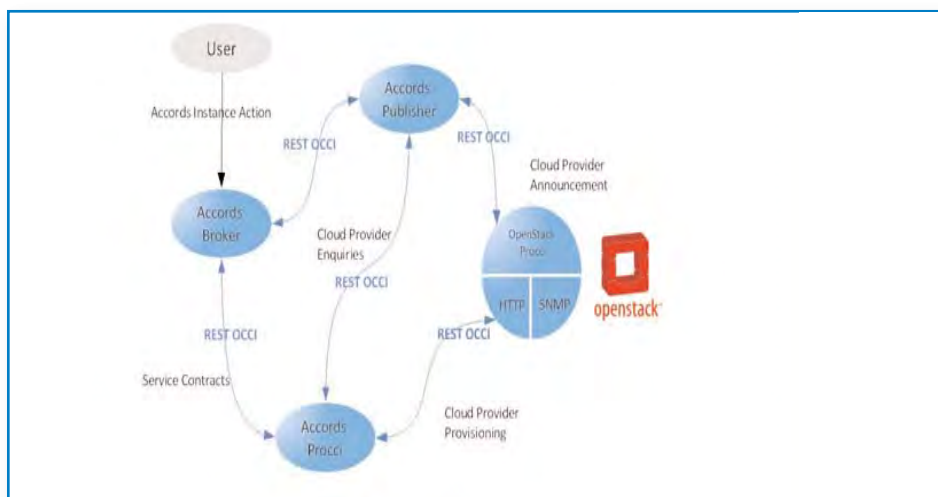


FIGURE 5. OPENSTACK PROVISIONING

OPENSTACK IMPLEMENTATION

Figure 5 shows how CompatibleOne uses an OpenStack platform to provision cloud resources:

- **THE USER SUBMITS THE MANIFEST** for parsing and subsequent brokering to deploy the required resources.
- **THE ACCORDS BROKER**, with the Publisher, issues a request for contract negotiation to the ACCORDS Procci.
- **THE PROCCI**, working with the ACCORDS platform placement tools, selects the most appropriate provider of the required type (in this case, OpenStack).
- **THE PROCCI NEGOTIATES** the technical aspects of the provisioning through the platform-specific Procci responsible for the dialogue with the OpenStack platform itself. This maintains a high degree of abstraction for the service modeling, with specialization handled over the last mile. The same approach is used to provision resources on the other platforms currently available through CompatibleOne—notably, OpenNebula*,

Windows Azure*, and Amazon EC2*—resulting in a homogeneous interface that makes it easy to aggregate and integrate heterogeneous resources.

It's easy to deploy a database running on one platform (e.g., OpenNebula) being used by a Web application running on an OpenStack platform, with interconnection performed during deployment by the Procci and CompatibleOne Software Appliance Configuration Services (Cosacs) components of the ACCORDS platform. The heart of the ACCORDS platform brokering system is the placement engine. This subsystem offers an extensible set of algorithms to select a provider type or platform operator, depending on the constraints and requirements specified in the input manifest.

Placement depends on the technical requirements for infrastructure and application software as well as on location and proximity, performance criteria, cost, and quality of (i.e., operator reliability). Because the placement engine is extensible by nature, almost any placement criteria will work. This could be a security enforcement constraint. For example, a provider's suitability might depend upon a proven level of

technical expertise in a particular security domain such as an OpenStack compute node protected with Intel® Trusted Execution Technology (Intel® TXT).

Energy efficiency is another important concern in cloud computing and data center management. The CompatibleOne Energy Efficient Services (COEES) component manages and processes energy consumption and efficiency information received from energy monitoring probes. The results and processing provide an important source of influence for the placement engine.

OPEN AND AGILE SOLUTION

CompatibleOne shows there is at least one solution to cloud computing problems like interoperability, portability, and reversibility. This simple yet powerful solution opens up new horizons and makes it possible to envision quick and agile development of new ideas and service concepts.

Cloud services can be complex because of the number and diversity of participants, the abundance and originality of their offerings, and the business model each one uses. Innovation and competition are the two key forces that govern the cloud market seg-

ment. But with solutions such as CompatibleOne, the marketplace—especially providers—will find themselves under increasing pressure from users to open their services up to meet these needs. This opening won't hinder their innovation, competitiveness, or competition. Instead, it will consolidate and strengthen the marketplace.

CompatibleOne, with ACCORDS and its PROCCI, makes it easy to link up offerings and make them interoperable. In the future, interoperability will come from automated negotiation of service contracts based on programmable SLAs. The programmable SLA will translate the clauses of a contract between the consumer and service providers into language that cloud services can interpret. It will then automatically negotiate service levels with providers through the contract and payment stages. This will make intermediation and automation of the contracting process easier—and also encourage new services.

CompatibleOne can serve as the foundation for an ecosystem that can include new value-added services, innovative start-ups, and new cloud providers.

To learn more about CompatibleOne, visit www.compatibleone.org.

Back to Contents

CLOUD DEPLOYMENT AND DESIGN

On Intel® Platforms

To bring root of trust into Intel's system on a chip (SoC) products and enhance security and integrity of the platforms, Intel has been working on a platform building block named processor secured storage (PSS).

MEETING INDUSTRY DEMANDS

This effort is aligned with industry demand for hardening the hardware and software integrity of devices and the global agenda of trust elevation, enabling secure machine-to-machine (M2M) transactions.

PSS is a fundamental technology building block comprised of a smart, secure, dual-port (I2C and RF) non-volatile memory (NVM) that is power-form-factor scalable for integration into the Intel® architecture compound (on package and eventually on die). It provides easy provisioning through the value chain and life stages of the platform, enabling Intel architecture and its other assets (i.e., firmware, middleware, OS, and applications) to store keys, certificates, and secure code onboard.

Uses include, but aren't limited to:


- **PROTECTION** against known and unknown firmware attacks and platform recovery
- **ENABLING** a variety of secure mobile services

- **ADDRESSING** identity management and multifactor authentication issues which are currently unresolved (and many other end user visible-valued secure services)
- **SECURE** M2M video
- **VOICE** and content
- **LOCATION-BASED** services

Today, there are a variety of platform storage solutions outside the Intel architecture. These include:

- **NON-VOLATILE** memories such as SPI flash
- **EEPROMS**
- **PROTECTED PARTITIONS** of mass storage
- **DISCRETE**, secure chips like SIM cards, near-field communications (NFC) plus secure elements, and trusted platform modules (TPMs)

Some of these are more trusted for storing keys and secrets; others are available for clear text data and code that may not require any high level of security. These solutions are available at different costs, form factors, and performance levels.



**THIS EFFORT IS
ALIGNED WITH
INDUSTRY
DEMAND FOR
HARDENING THE
HARDWARE AND
SOFTWARE
INTEGRITY OF
DEVICES AND
THE GLOBAL
AGENDA OF
TRUST
ELEVATION.**

BASIC PLATFORM DEFINITIONS

PSS is a very low-power, cost-effective, and small-form-factor dual-ported, immutable NVM. PSS is based on the UHF EPC Global Gen2 RFID IC* product with DC input and a I2C interface. PSS provides an easy provisioning capability via radio frequency or I2C through the value chain of the platform, enabling Intel architecture and other platform assets (firmware, middleware, OS, and applications) to store desired tokens, certificates, and secure code onboard.

The RF front end of the PSS provides two RF input ports, enabling both near- and far-field antenna design, which can be individually enabled or disabled, or even permanently shut off, depending on design requirements.

PSS is available in versions with 2.1 Kbits and 8 Kbits of user space, with either four or 15 one-time programmable (OTP) memory blocks for storing immutable keys and tokens using a I2C interface.

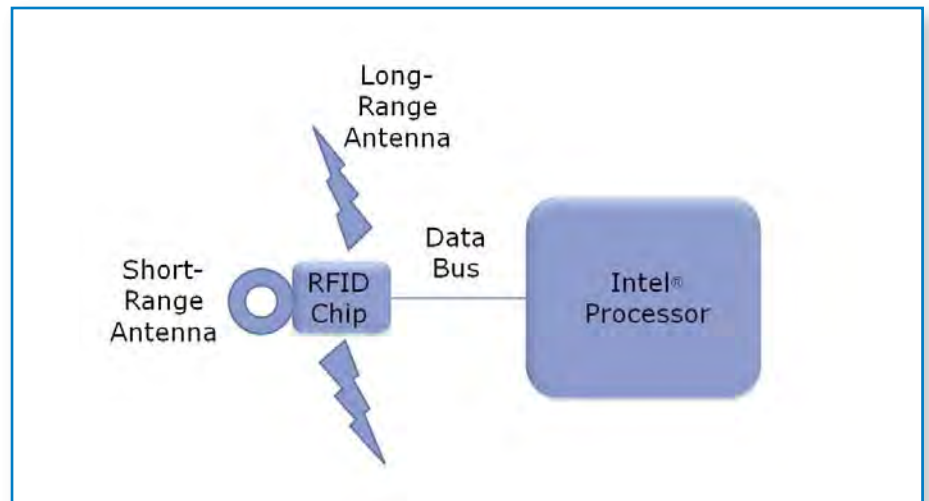


FIGURE 1. PROCESSOR-SECURED STORAGE

BASIC PLATFORM DEFINITIONS

TRUSTED EXECUTION ENVIRONMENT

The term trusted execution environment refers interchangeably to current or future secure engine implementations of Intel SoCs such as secure engine coprocessor (SEC). In general, TEE refers to the integrated hardware security engine and associated firmware isolating and managing various assets of the platform such as the processor secured storage.

SECURE ENGINE COPROCESSOR

SEC is one of the trusted execution environment implementations targeted for the current generation of Intel processor-based tablets and embedded devices.

PROCESSOR-SECURED STORAGE

Processor-secured storage technology includes a secured NVM managed by the TEE. This technology enables a variety of secure mobile services and addresses both unresolved identity management issues and many other end-user-visible and valued secure services. Features include:

- **DC (I2C) AND RF** dual interface
- **EPCGLOBAL UHF** Gen2 RFID air interface

SECURITY FEATURES

- **A UNIQUE 96-BIT ID** and a total of between 2.1 Kbits and 8 Kbits of user memory space
- **A DISTINCT NUMBER** of one-time programmable banks of NVM via RF and I2C (16 for the 8-Kbit version)
- **PASSWORD-BASED NVM** write control on RF link and QTI for read control and data privacy on RF link
- **RF-I2C ARBITRATION** with programmable auto RF disable

PSS software stack support includes:

- **EPC GLOBAL RFID** software stack and applications available
- **I2C USERLIB** and driver support for Windows* and Android*
- **JAVA APPS UVM* INTERFACE** (available from Intel under non-disclosure agreement)

Verified Boot is a hardware-based verification feature of the firmware images that is mandatory for the Windows 8 platform. This capability is complementary to the Windows 8 secure boot. The PSS ROM storage managed, by the TEE for storing keys, is more secure and is the Intel

preferred option in Windows 8.

Note that Verified Boot uses Intel-generated keys fused in the SOC to verify the firmware components for the secure boot. Generated by Intel, these keys are distributed to all OEMs and ODMs. With PSS, OEMs can use OEM-specific keys instead of Intel-generated keys to verify the first stage of BIOS.

This enables OEMs to ensure that a specific OEM's BIOS works only on each specific OEM's hardware, ensuring a unique provisioning solution supported by Intel's firmware.

In Windows 8 secure boot, the security trust starts from Intel architecture firmware (BIOS) after the Intel architecture core powers on. In Intel Verified Boot (with PSS), the root of trust starts from the hardware power-on (Intel's Verified Boot and Windows 8 secure boot are complementary and both necessary).

Besides the Verified Boot model, there are other usage models to ensure user identity and trust for

enterprise tablets as part of banded platforms using processor-secured storage for required tokens.

BANDED PLATFORM TRUST ELEVATION THROUGH PROCESSOR-SECURED STORAGE

As part of multifactor authentication, a banded tablet can be connected to the private cloud only if the provisioned platform ID and biometrics match beyond the traditional user login and password.

Additional contextual credentials can also be used (e.g., GPS coordinates and last login plus associated policies set). In this scenario, three core ingredients are assumed present:

- **AN END-POINT DEVICE** with a unique platform identity registered/provisioned for registered user(s) in an immutable storage (i.e., processor-secured storage)
- **SOME RELIABLE IDENTITY FRAMEWORK** (e.g., a biometric method such as iris, fingerprint, or keystroke signature to elevate trust)

ACCESS AND PLATFORM

- **A CLOUD AUTHENTICATION**

method combining the user ID and password, the platform identity, and biometric credentials

Access is granted to user(s) and, depending on multi-factor authentication, the user can perform basic or restricted actions including secure peer-to-peer (P2P) communication (e.g., secure P2P video, voice over IP [VoIP], or file sharing).

BANDED PLATFORM

A service provider private cloud (e.g., an enterprise IT cloud) that has determined its electronic authentication requirement at NIST Level 3 or higher manages its banded platforms. There are number of required ingredients to establish trust to such a private cloud including three out of these four typical, basic elements required for a verified session:

- **WHAT YOU KNOW:** Shared secrets such as logins, passwords, or other public/private information
- **WHO YOU ARE:** Secured biomet-

rics coupled with device identity including liveliness

- **WHAT YOU HAVE:** Hard or soft tokens, immutable device identity, or an alias
- **CONTEXTUAL INTELLIGENCE:** Location, time, or chronology of events

The service provider receives an electronic identity credential from an end user that is recognized as a Level 1 credential (login plus password). By applying one or more recognized methods for assessing the identity of the end user, the service provider can ensure that the presented credential actually represents the asserted identity at higher level(s) of assurance comparable to NIST level 2, 3 or 4.

NIST levels 3 and 4 can be achieved by electronic product identity provisioned into the banded platform by IT and stored encrypted in the end-point device's secured storage/vault. This vault/secured storage is as atomically close to the SoC as possible. Both initially and dynamically, it is pulled by the cloud server for verification.

Cloud and data center policies are applied and managed as required, depending on various trust level demands. For example, this could include biometric or encrypted GPS credentials or an IP address, using the white list provisioned by the cloud login server as an auxiliary credential for multi-factor authentication.

PEER-TO-PEER FLOW

Once the tablet is identified as a banded platform of the private cloud (Figure 2), the cloud creates a shared key and stores it in the dedicated region of processor-secured storage of targeted devices requesting P2P connection (Figure 3).

As a function of the policies provisioned, the server dynamically queries the device. If all credentials remain intact, it maintains the connection. If not, it terminates the session. (Note that we may use GPS or IP as second-level credentials. If the Level 1 credential does not match the platform provisioned for the user, a kill pill is issued for the device and the associated login name.)

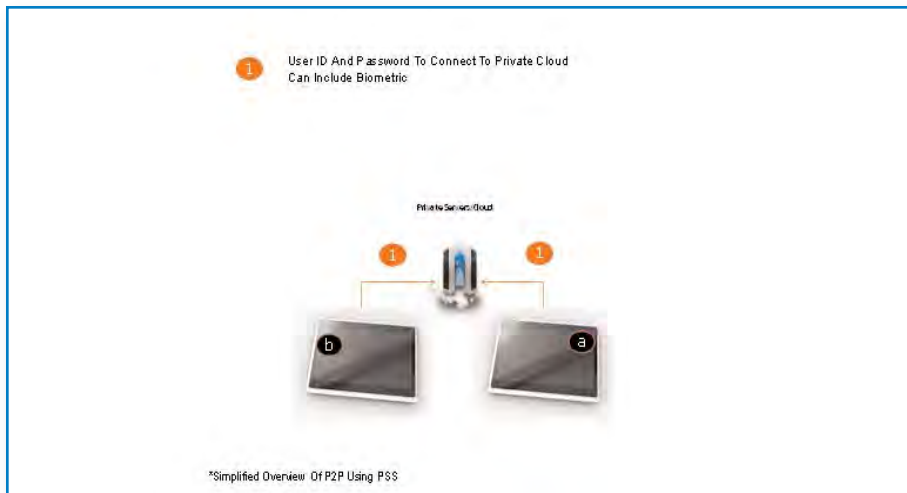


FIGURE 2. TABLET ID AS A BANDED PLATFORM OF THE PRIVATE CLOUD



FIGURE 3. THE CLOUD CREATES A SHARED KEY

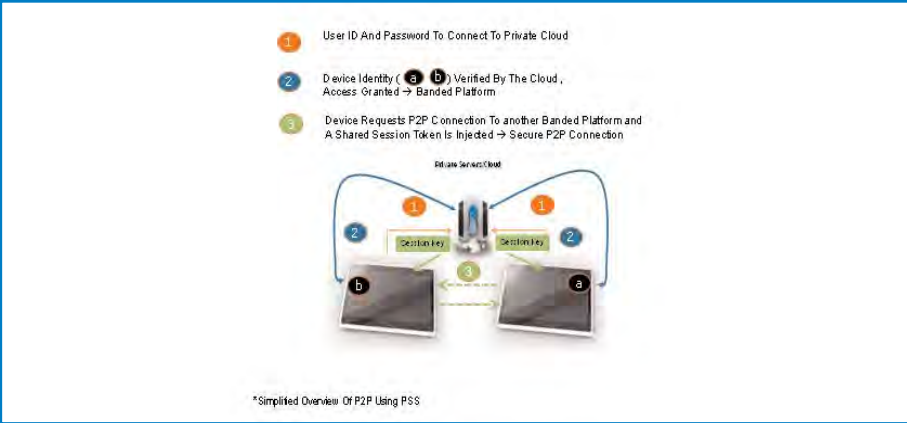


FIGURE 4. CROSS-DEVICE CONNECTION

Once the client-to-cloud connection is created, we can create a cross-device connection enabling secure video, content sharing, and VoIP (Figure 4). A number of solutions are under investigation for identity assurance.

CLOUDWIDE IDENTITY ASSURANCE FRAMEWORK

Intel and aTrust (www.atrust.ca) have collaborated on an identity assurance framework inspired by both Identity 2.0 (also called digital identity), a set of methods for identity verification on the Internet using technologies such as information cards or OpenID, and the personal identity framework.

Called Cloudwide Identity Assurance Framework, this new framework unambiguously assures and authenticates an individual every time they a subscribing service provider asks for authentication. The consumer’s biometric authentication is bound to their bona fide civil credential(s) through the consumer’s personal, Intel-powered trusted computing platform and aTrust’s assurance, authentication, and security mechanisms.

Intel and aTrust developed the Cloudwide Identity Assurance Framework to address the alarming rise in identity theft and electronic fraud as well as to fulfill the promise of best identity and authentication practices. aTrust's Identity Assurance Framework*, augmented by the capabilities of Intel® Trusted Execution Technology (Intel® TXT), has produced a solution that is both business-focused and consumer-centric, addressing today's increasing e-business risks and challenges. The result is a powerful and unique infrastructure that gives both consumers and service providers the assurance they need to transact e-business routinely and in confidence. This next-generation identity system provides participating service providers with a high degree of certainty as to the true identity of the individual seeking online access. Figures 5 through 7 show the key computational elements and communication paths among the elements. Participating service providers subscribe to aTrust's Cloudwide Identity Assurance Service*, installing aTrust's Service Provider Module*,

which exposes an aTrust identity assurance application programming interface (API) to the service provider's Web applications and supports secure communications with the aTrust Cloudwide Identity Assurance Service and the consumer's biometrically-embedded trusted computing platform.

A key subsystem, from the consumer's perspective, is the trusted computing platform and software that is produced, integrated, and initialized in volume by Intel, aTrust, and others. The initial versions will be deployed on tabled PCs with integrated biometric hardware. Consumers will buy these systems from qualified retail computer, software, and mobility dealers.

After acquiring this computing platform, a consumer enrolls biometric information and creates a PIN number supported by aTrust Cloudwide's embedded Identity Assurance and Security Module. The consumer then subscribes to aTrust's Cloudwide Identity Assurance Service* and initializes a personal profile that includes an identity supported by a photo and other public and private infor-

mation. This information lives in private data stores on the computing platform.

The consumer then seeks out a registration authority (e.g., a local department of motor vehicles or passport office) and has the credential(s) verified in person. The agent validates the consumer's biometric identity and tags it online through an identity assurance service.

After completing these enrollment steps, the consumer is ready to conduct e-business with any aTrust subscribing, cloud-based service provider or with a subscribing enterprise with a private cloud.

On the consumer's first visit (and attempted access) to the application, and whenever required thereafter, the service provider requests the Cloudwide Identity Assurance Service to provide an assurance token for the consumer, validating identity. Next, the service provider obtains the consumer's public key as well as other vital information including a digital certificate used to support private and secure communications with the consumer.

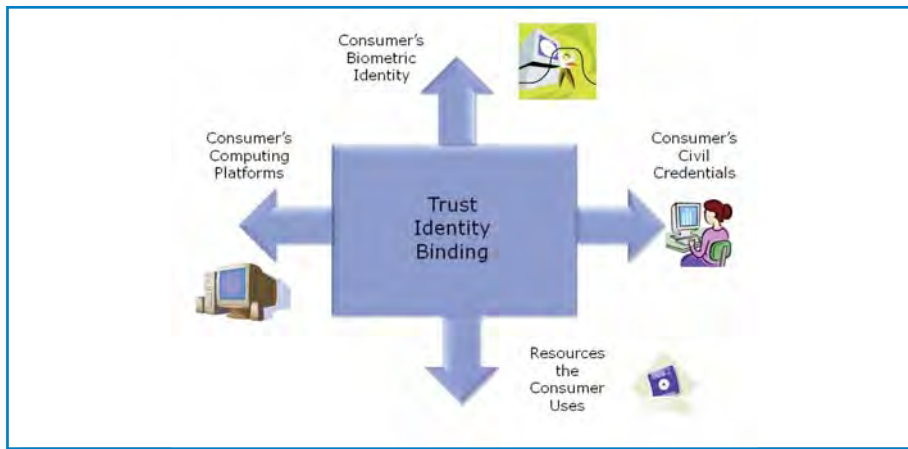


FIGURE 5. IDENTITY BINDING AND ASSURANCE

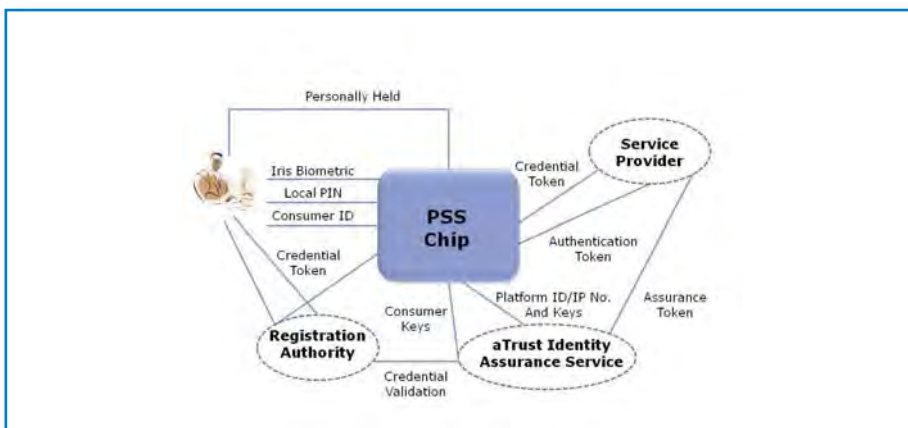


FIGURE 6. IAF IDENTITY INTEGRATION UNDER IRIS

Note that the service provider is getting information to qualify the consumer for possible enrollment without having met the consumer in person. To complete the sign-in enrollment process, the service provider can also request additional information and electronic credentials from the consumer before deciding to grant access to the Web application.

Once the user is enrolled, the service provider asks for an authentication token,

using the consumer's digital certificate to verify the authentication token before granting access to the Web application.

During consumer sessions, the service provider can request reauthentication of the consumer (e.g., for high-value or risky transactions) at any time. The Cloudwide Identity Assurance Service can also send consumer problem notices (e.g., consumer subscription expiration or changes) to the service provider, which may prompt the

service provider to take some remedial action.

Once a consumer's digital identity has been verified in person, he or she can engage in secure, identity-verified e-commerce from a remote location over the Web or on an enterprise network with a much smaller risk of identity loss and electronic fraud.

Both the consumer and service provider benefit from the reduced cost of doing business that comes from single-sourcing the in-person credential verification processes and using biometric authentication for a single sign-on.

The Cloudwide Identity Assurance Framework uses Intel's TEE, its embedded PSS chip, and the Beihai Virtual Machine*, all of which are enabling technologies for the development of secure identity solutions.

The PSS chip is embedded with a unique platform/device identifier and private encryption key, which together enable aTrust's identity assurance mechanisms and services to authenticate any Intel platform, securely binding aTrust's assurance services and Intel's trusted computing platform.

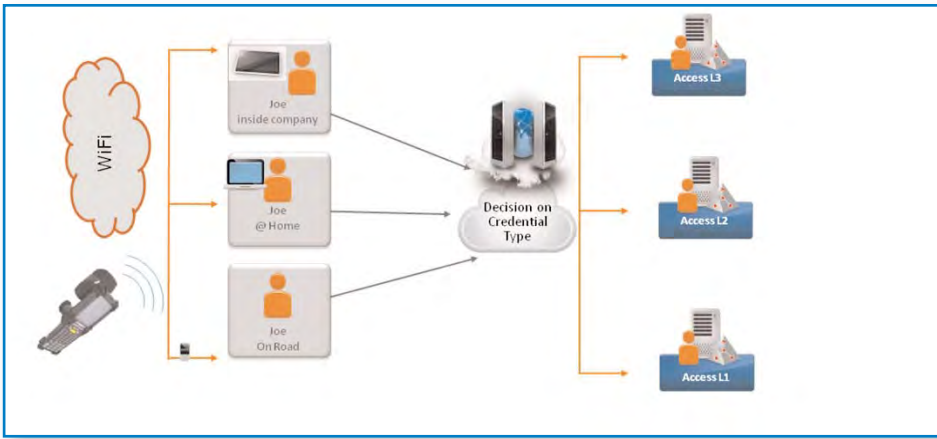


FIGURE 7. DECISION ON CREDENTIAL TYPE

The Intel platform also supports iris biometrics, which execute with the Beihai protection domain using aTrust authentication and security mechanisms that positively identify the individual consumer using that platform.

Biometric templates (minutia) are cryptographically protected within the Intel platform's protection domain.

aTrust's identity assurance modules and services generate private keys for the consumer, storing them, and other private information such as digital certificates and electronic credentials, in the PSS chip.

aTrust's Identity Assurance Framework supports in-person validations of the civil identity creden-

tials of consumers and uses Intel's platform to biometrically bind consumers to their civil credentials.

Such capabilities enable enhanced consumer identity assurances. Authentication greatly reduces identity administration and management costs for online consumers of the service provider.

IDENTITY ASSURANCE FRAMEWORK INGREDIENTS

The Identity Assurance Framework uses several mechanisms to implement identity binding and give service providers the identity assurances they need (Figure 6):

- **THE CONSUMER'S COMPUTING PLATFORM** is a trusted appara-

tus that is cryptographically bound to aTrust's Cloudwide Identity Assurance Service when deployed.

- **EACH CONSUMER** is biometrically bound to their asserted identity by Intel's computing platform and by an aTrust Consumer Module.
- **EACH CONSUMER AND THEIR ASSERTED IDENTITY** is cryptographically bound by the aTrust Consumer Module to the identity assurance service.
- **THE CIVIL IDENTITY CREDENTIALS** of consumers are verified and computationally bound to each consumer's identity by aTrust's Cloudwide Identity Assurance Service.
- **CRYPTOGRAPHICALLY PROTECTED TOKENS** are provided by aTrust to service providers on request, providing assurances of the consumer's civil identity.
- **CRYPTOGRAPHICALLY PROTECTED TOKENS** provided by the consumer to the service provider communicate the biometrically-authenticated identity of the consumer.

- **THE CONSUMER** can elect to release cryptographically-protected civil identity credential identification numbers and other private information stored in their Private Data Store to the service provider.

LOCATION-BASED ACCESS CONTROL AND SERVICES

The three factors (who you are, what you know, and what you have as contextual data such as location) can be injected into the processor-secured storage OTA, which can determine the access control. In this scenario, access to different classes of data is permitted only if the right token

for the associated location is programmed into the PSS (OTA or via RF as associates to enter or exit the boundaries).

UNAMBIGUOUS IDENTITY PROTECTION

In summary, Cloudwide Identity Assurance Service is uniquely capable of delivering unambiguous assurance of a consumer's actual identity every time they are asked to be authenticated by a subscribing service provider. This is made possible by integrating in-person verification by a registration authority using aTrust's identity assurance service, with biometric authentication of the consumer on a personally-held trusted computing

platform provisioned by Intel.

This approach enables aTrust to provide assurances to service providers as to the real identity of consumers without observing, storing, or releasing their credentials to service providers. The consumer remains in total control over their private information, including credential details, which they elect to store on their trusted computing platform and which they may choose to release to other parties. aTrust's Cloudwide Identity Assurance Framework brings into focus the consumer's true personal identity—namely, their digital identity.

To learn more, visit www.atrust.com

[Back to Contents](#)

IN AN UNCLEAR HPC CLOUD LANDSCAPE, AN EFFICIENT HPC-ON-DEMAND SOLUTION Bull's extreme factory*

Olivier David,
ISV Alliances Manager, Extreme Computing Business Unit, Bull
olivier.david@bull.net

High-performance computing (HPC) as a service is certainly not a new concept. It's been available under other names (e.g., application service provider or ASP, hosting, on-demand, grid computing), in various forms, for the last 20 years. Some of these concepts have been moderately successful. Others have failed miserably.

THE CHANGING LANDSCAPE

Today's landscape has changed—technically, financially, and culturally.

Technically, wide-area network (WAN) and Internet network bandwidth have increased tremendously, making it much easier to transfer important data sets. Virtualization, though not as crucial in HPC environments as in the traditional IT space, has created many possibilities for sharing and multi-tenant infrastructures. Financially, both hardware infrastructures and network costs have gone down while the costs for people, software, and space have stayed relatively stable.

But the biggest change has actually occurred in the last five or six years with cloud computing being hyped, buzzed about, feared, over-emphasized and, finally, actually adopted by many successful companies. Business models have been validated, successes registered, and now the concept has turned into a reality.

HPC is following the same path. Companies still worry about letting

their data outside their premises. They still have network bandwidth issues with their ever-increasing data sizes. They still have to negotiate with their internal IT and finance departments. But increasingly, they're overcoming these obstacles and making the last step towards using real HPC as a service solutions.

The software as a service (SaaS) market segment is exploding. HPC as a service, still a small part of the market segment today, is quickly gaining momentum. New solutions are available every month, with prices dropping regularly. Companies of all sizes and industries are looking at solutions, carrying out proofs of concepts, and confirming orders for both public and private cloud offerings.

With a clear vision of company needs and marketplace expectations, and using its experience from similar projects, Bull has designed a new offering called extreme factory to cover HPC as a service.



**INCREASINGLY,
COMPANIES ARE
MAKING THE LAST
STEP TOWARDS
USING REAL HPC-
AS-A-SERVICE
SOLUTIONS**

FLEXIBLE, ON-DEMAND SOLUTION

This flexible, on-demand HPC offering is for companies without enough compute resources to satisfy their HPC workloads (Figure 1). With extreme factory, organizations of all sizes can innovate without making major investments in powerful computing resources. All a company needs to use the solution is Internet access via a dedicated portal. Bull provides the infrastructure to run workloads in targeted turn-around times. Users pay for the time used, enabling them to adjust operating costs to the schedules and goals of each project. The extreme factory cluster is hosted at the Bull data center in Les Clayes-sous-Bois, France (near Paris), in a highly secure environment.

extreme factory is a complete offering, operated 100 percent by Bull and its subsidiaries. Bull's experience and expertise in systems design, services operations, applications management, Web development, security components, and telecom-




FIGURE 1. BULL'S HPC CLOUD COMPUTING SOLUTION

munications was key to bringing to the marketplace a fully functional and flexible solution.

User needs vary, but Bull has implemented three usage models with the flexibility to meet different requirements:

- **DEDICATED:** Resources dedicated to users with a long-term need (six months or more). Bull allocates dedicated hardware resources (i.e., compute and service nodes, storage) to ensure guaranteed stability and security. This model can be easily customized, and either the user's operator or Bull can add optional VPNs.
- **RESERVED:** Guaranteed resources through reservation. Many users have peak loads requiring added resources. This model allocates compute time in periods of one or more weeks by dedicating virtual login nodes and provisioning the associated physical compute nodes for the duration requested.
- **SHARED:** Resources are mutualized and allocated on a first in, first out basis. This model is closest to the traditional commercial cloud model.

HIGHLY SECURE ENVIRONMENT

Users need to carefully consider the business impact of these three usage models before deciding which one to choose.

Both the dedicated and reserved models are billed when the user makes a reservation, matching user needs to planned projects. Even with flexibility built into the model, the user has overall responsibility for making use of the time paid for. The shared model is more like the pay-as-you-go model for traditional clouds. Users buy time credits in advance. Bull sends email warnings when credits are running low. With this model, the user does not pay for unused resources.

extreme factory also provides solutions for companies in key vertical industries including manufacturing, financial services, healthcare, life and material sciences, media, and oil and gas.

extreme factory is especially well suited to companies in the manufac-

turing market segment, which have used all of the solution's major crash and computational fluid dynamics (CFD) applications. These applications scale well, and companies have used them to parallelize dozens or even hundreds of cores. For example, L&L Products, a tier 1 subcontractor for major automotive OEMs, uses extreme factory's PAM-CRASH* code, coupled with Digimat*, for modeling and assessing light fiber products, replacing heavy traditional metal designs. (Learn more [here](#).)

Some ISVs are using extreme factory to gain new customers licensing models dedicated to on-demand uses. CD adapco, for example, has launched its Power-On-Demand* service to give users flexibility for optimizing license usage.

In life sciences, companies use extreme factory to run genomics simulations. Other developing market segments include:

- **MEDIA:** Film rendering typically uses hundreds of cores for weeks or months, making it impractical to buy and operate the necessary hardware for such restricted durations
- **OIL AND GAS:** Seismic and electromagnetic data needs to be processed with fast turnaround times to save millions of dollars in drilling campaign decisions.

Figures 2 and 3 show how companies are using extreme factory in scientific applications.

extreme factory infrastructure is a highly secure environment that is physically accessible only to authorized Bull operators. The cluster is a complete HPC system, with state-of-art CPUs and high-speed InfiniBand Interconnect*.

Two types of compute nodes are available in the extreme factory supercomputer and regularly updated to accommodate state-of-art technology:

HIGH SPEED, LOW LATENCY

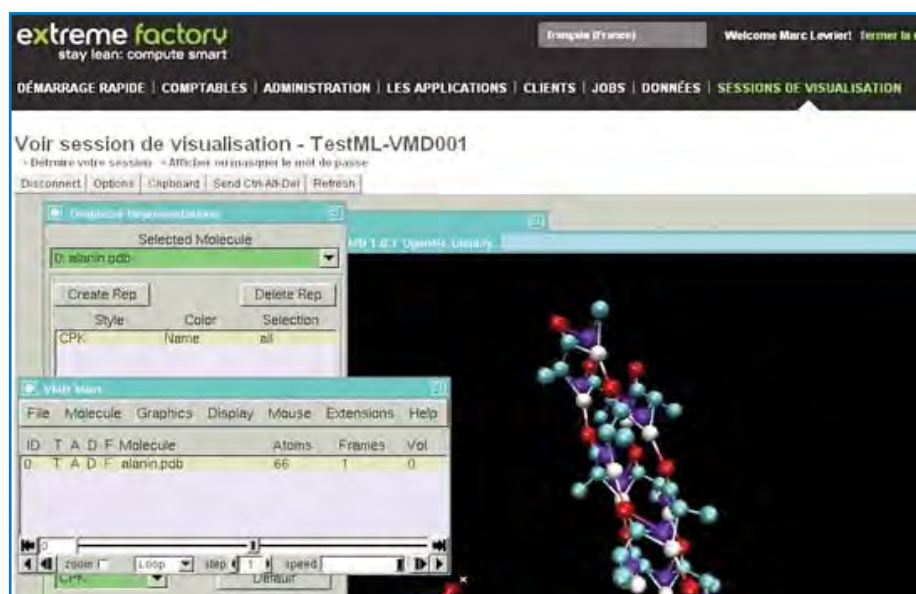


FIGURE 2. USER INTERFACE

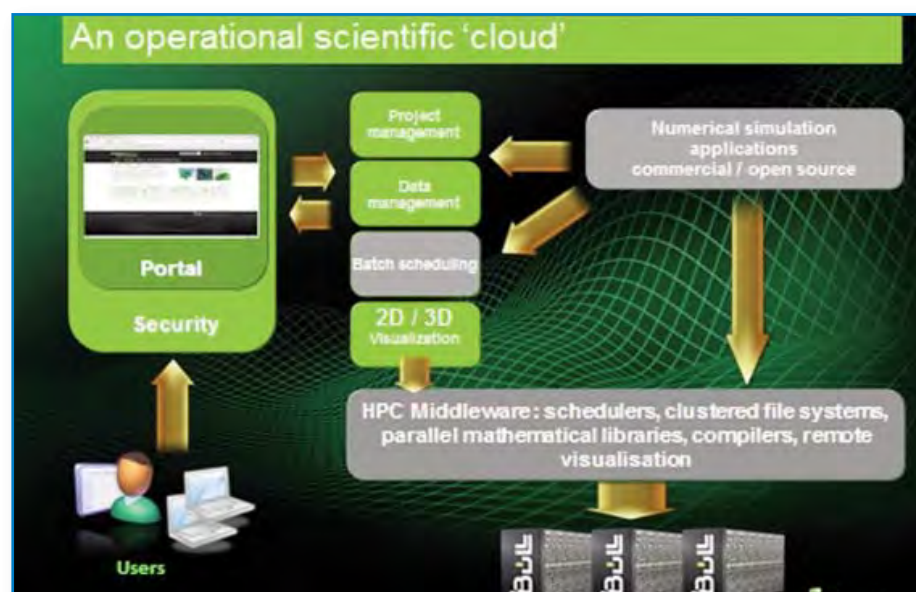


FIGURE 3. A MODEL CLOUD

- **BULLX* B500 AND B505**
BLADES have highly-optimized Intel® Xeon® processors 5600 series with 24 gigabytes of memory. bullx B505

blades have additional dual Nvidia M2070* GPUs. They are currently being transitioned to bullx B510 blades using latest Intel Xeon

processor E5-2600. These blades have already been deployed by Bull in 2 petaflop supercomputers (IFERC* in Japan and GENCI* in France) and are ideally suited for all HPC loads.

- **BULLX R423-E2 OR BULLX S6030** large memory nodes are available for specific operations like preprocessing data models, which require memory configurations of 512 gigabytes or more. Service nodes are either physical bullx R423-E2 nodes with good I/O and memory capabilities or virtual nodes for login and security isolation and R425-E2 nodes for visualization.

All nodes run a standard Linux* environment based in RHEL* 5 or 6 compatible kernels as well as Bull's own cluster management software. Windows* nodes can also be deployed for individual projects. All compute nodes are connected through a high-speed, low-latency quad data rate (QDR) InfiniBand network, which is necessary to achieve near-linear scalability of most applications.

WELCOME TO THE FUTURE

There are two types of storage:

- **A HIGH-PERFORMANCE** parallel Panasas* storage cluster
- **A DEDICATED, HIGH-CAPACITY** NetApp* system for long-term storage

Standard communication needs are addressed by secured lines with bandwidths of 100 megabytes to 1 gigabyte per second. These are accessible through the Internet with optional VPNs. For companies with higher bandwidth requirements, Bull PI, a Bull subsidiary, advises installing adequate point-to-point links at speeds up to 10 gigabits per second. This configuration is very flexible and scalable to accommodate CFD and crash applications, which can easily scale to hundreds of cores. Users can run these jobs, although often with at a lower number of cores per job. The total average machine size is approximately 150 teraflops, with

the exact size easily adjustable to the current workload. No company has ever saturated the infrastructure, although requests for 800 or more cores are common.

Interestingly, companies do not often request high availability, a useful traditional HPC feature, in extreme factory. This feature is available only in a dedicated solution with custom environments. It requires a high quality of service (QoS) on a supercomputer where all administrative tasks are centralized and controlled by expert Bull administrators.

Security is integrated part of the extreme factory architecture, not an afterthought. In all configurations in physical or virtual modes, customers are fully isolated with full and unique access to their data, projects, and jobs. They cannot access any other information, and don't even know

which other customers are using extreme factory, nor which applications are being used. The preferred mode of interaction is https. Also possible is ssh, with restrictions so that security is not compromised.

EXTREME FACTORY SOLUTION

Extreme factory was built with 10 years of experience from Bull's people and subsidiaries to meet companies' needs for on-demand HPC. Bull has built a complete solution that offers a full worldwide infrastructure to run HPC jobs efficiently and securely.

Bull is also planning to introduce a private cloud version of extreme factory, which it is currently testing with customers. Operating inside the customer's WAN, it can burst extra requests to the public extreme factory version. Bull believes this is the future of on-demand HPC service.

To learn more about extreme factory, visit www.extremefactory.com.

[Back to Contents](#)

FUTURE HOSPITAL

A male doctor with dark hair, wearing a white lab coat over a light blue shirt and a patterned tie, is holding a black tablet. He is looking down at the tablet with a slight smile. A male patient with short dark hair, wearing an orange V-neck sweater over a white collared shirt, is looking up at the tablet. The background is a clinical setting with a light blue wall and a framed anatomical chart of the human skeleton.

Cloud computing is a central topic in the IT industry. But in healthcare IT, the common perception is, "Cloud computing is interesting, but it's not truly suitable for healthcare." After turning to topics like keeping patient data secure, most cloud computing discussions end fairly quickly.

CLOUD COMPUTING IN HEALTHCARE?

Intel and German healthcare provider Asklepios Clinics wanted to find a way for organizations to enjoy cloud benefits like resource efficiency, scalability, and automation by adapting cloud computing to the specific demands of healthcare. The two companies have developed a distributed health cloud architecture as part of the Asklepios Future Hospital* (AFH*) Program.

Established in 2006 by Intel and Microsoft, the AFH Program represents an alliance for the future of healthcare, connecting roughly two dozen leading international companies to work together on innovative solutions for healthcare and prepare for the challenges of tomorrow.

An important part of the AFG Program is the Distributed Health Cloud Project, a collaboration between Intel and Asklepios. With its highly distributed health facilities, the Asklepios Group's IT requirements favor a distributed, high-performance server platform that can efficiently scale to support different-sized locations. The automation intro-


duced by cloud computing fits Asklepios' requirements for localized self-service and will give the group the freedom to bring its IT to the next level.

CONTROLLING COSTS

The cost of medical care has spiraled over the last decade, primarily because of population growth in western countries and ongoing innovations in medical treatment. In contrast, human and financial resources are limited. Efficient use of these resources is essential, especially in healthcare.

Asklepios has already proven with its OnelT* concept that a highly standardized IT infrastructure can help meet efficiency demands.

Now, using cloud computing technologies in healthcare, Asklepios is looking to address more than just efficiency, since the availability and agility of IT services are essential to clinicians, even in the smallest facility. The redundant, multi-node nature of cloud platforms addresses this concern. Through automation and infrastructure as a



CLOUD COMPUTING CHALLENGES

service (IaaS)-type self-service, hospital IT personnel can quickly and easily adapt virtual computing resources and meet the growing demands of doctors and nurses for IT throughput without having to waste time and money acquiring basic physical IT infrastructure.

CLOUD COMPUTING CHALLENGES

At Asklepios clinics, the major challenge for cloud computing is meeting the demand for local data and responsibility.

This demand is based on regulatory requirements and technical limitations. On the regulatory side, the processing of patient data is regulated by federal hospital laws as well as regulations in each state. In three states, for example, patient data cannot be processed by third parties. German authorities demand physical or logical separation of patient data of different hospitals and practically prohibit a transfer of health data to and from non-

European countries. This limits the potential benefits of “locationless” cloud computing (e.g., international support, a master patient index, and easy visibility of data to different hospitals and medical experts). On the other hand third-party processing by company data centers or national cloud providers is allowed if contracts and technical controls such as end-to-end encryption effectively leave the hospital in control of its data.

On the technical side, it's nearly impossible to cost-effectively meet medical applications' steep demands for bandwidth, latency, and availability with current WAN technologies.

This is especially true for small healthcare facilities in rural areas.

The reasonable demand for high availability of critical IT services for the patient treatment process (e.g., the lab information system) leads to a complex and cost-inefficient solution when aiming at a central cloud.



With a market segment share of more than 20 percent, the Asklepios Group is one of the three largest operators of private hospitals in Germany. The group's strategy—which focuses on high-quality, innovation, and sustainable growth—has been rewarded with dynamic growth since its formation in 1984. With its more than 44,000 employees, Asklepios runs more than 140 health facilities.

For more information, visit www.asklepios.com.

DISTRIBUTED HEALTH CLOUD

THE CONCEPT

The constraints of the healthcare environment—many local hospitals with application demands, data locality requirements, and limited WAN connectivity to Asklepios' central IT site—led Asklepios and Intel to take an unusual architectural approach. The team decided to trade off some resource pooling efficiency of the targeted IaaS platform for the capability to keep data local to each hospital and benefit from good local network performance. The clear separation between operating the IaaS platform from a central IT point of control and enabling local IT specialists to maintain control over installed workloads had to be kept intact. Ultimately, the team reached this goal with a highly standardized setup that included:

- **A CENTRALIZED AUTOMATION LAYER** based on Microsoft System Center* hosted at Asklepios' main IT site. This layer triggers VM management actions on each hospital's local virtualization platform (a pool of

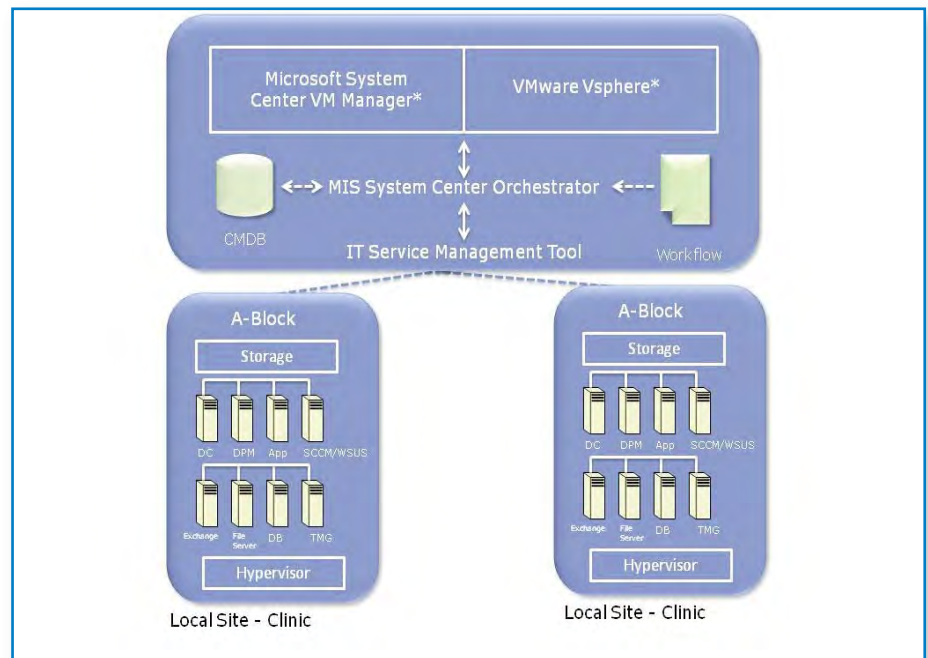


FIGURE 1. ASKLEPIOS HEALTH CLOUD STRUCTURE

VMware* or Hyper-V* hosts) over the WAN connection.

- **ASKLEPIOS' CENTRAL IT SERVICE MANAGEMENT TOOL** serves as the main entry point for local IT tasks, triggering orchestration workflows on this automation layer (Figure 3).
- **SINCE THE PLATFORM SPECIFICALLY RESTRICTS VM MIGRATION** between sites, it effectively turns the data locality constraint into the benefit of greatly reduced WAN bandwidth and latency requirements.

Figure 1 gives a schematic overview of the Asklepios Health Cloud components.

THE A-BLOCK: FLEXIBLE CLOUD POD Since local hospitals have different application performance requirements and different numbers of staff and patients to serve, the team conducted an analysis of currently running workloads. The study showed capacity demand across Asklepios' hospitals falls into several distinct throughput and redundancy categories.

PLATFORM CONFIGURATIONS

This clustering of requirements let the team create a specification for a hardware platform architecture that is both fixed in terms of functionality but scalable in terms of performance. Table 1 shows the three standard platform sizes that only differ in the number of virtualization hosts and usable throughput and internal compute redundancy.

These functions are identical for all configurations:

- **HARDWARE** is pre-integrated in a standalone rack with integrated HVAC, fire suppression, and access control.
- **A TOP-OF-RACK SWITCH** with 48 1GbE ports and dual 10GbE uplink provides the platform-local switching infrastructure.
- **A BOTTOM-OF-RACK**, integrat-

ed, uninterruptible power supply (UPS) provides battery backup in case of main failure.

- **AN INTEGRATED STORAGE ARRAY** (2U NetApp* FAS3240 plus 6U drives chassis chosen for the proof-of-concept setup) implements the centralized, persistent storage pool. This device was chosen because of the option to use its storage-level remote replication features for a site-redundant expansion of the current architecture should disaster tolerance requirements so dictate.
- **AN ISCSI-BASED TAPE DRIVE** for remote backup is associated with each A-Block but not integrated into the frame.

The main properties of the three platform configurations listed in Table 1 are based on these arguments:

- **A SINGLE VIRTUALIZATION HOST** based on two Intel Xeon processors E5-2640 with 192 GB of RAM (out of 24 x 8 GB DIMMs) represents the maximum memory configuration with the lowest cost per gigabyte.
- **ACCORDING TO WWW.SPEC.ORG**, a highest score of 444 could be achieved for the cited processors at the time of publication on the SPECint_rate_base 2006 benchmark.
- **DUE TO AVAILABILITY REQUIREMENTS**, the platform must tolerate the failure of a single virtualization host. Since a dual-host failure tolerance is not required, the minimum possible redundancy—effectively always the memory and throughput capacity of a single host reserved as spare—is used for each configuration.

TABLE 1. A-BLOCK SIZES BASED ON VIRTUALIZATION HOSTS WITH INTEL® XEON® PROCESSORS E5-2640

Configuration	No. Hosts	Raw Memory [GB]	Raw Throughput [SPECint_rate_base 2006 x No. Hosts]	Redundancy	Net Memory [GB]	Net Throughput [SPECint_rate_base 2006]	Target Host Utilization
S	2	384	888	1+1	192	444	50%
M	4	768	1,776	3+1	576	1,332	75%
L	6	1152	2,664	5+1	960	2,220	83%

- **THE NET VALUES FOR TOTAL RAM** and total throughput are calculated from the accumulated SPECint_rate_base 2006 score reduced by the employed redundancy.
- When distributing workloads equally among all hosts, it's essential to obey a memory and CPU utilization limit to keep available enough total spare capacity in line with the desired redundancy.

PROJECT STATUS AND OUTLOOK

Asklepios and Intel are creating a full-scale proof of concept setup of an A-Block for

demonstration and performance analysis. The central portal is already implemented. Figure 3 shows the tasks it supports for hospital-resident local IT personnel.

Asklepios corporate IT plans to use the A-Block concept as the foundation for its hardware platform standardization. On the procurement side, it hopes to reach economies of scale by reusing identical components and making effective use of local knowledge. Since it is easy to replace the virtualization host component with latest-generation platforms, the platform will be able to support performance demands into the future.



FIGURE 2. S-SIZE A-BLOCK (SERVER WILL BE REPLACED BY THE UNITS FOR BETTER SCALABILITY)

To learn more, visit www.asklepios.com

[Back to Contents](#)

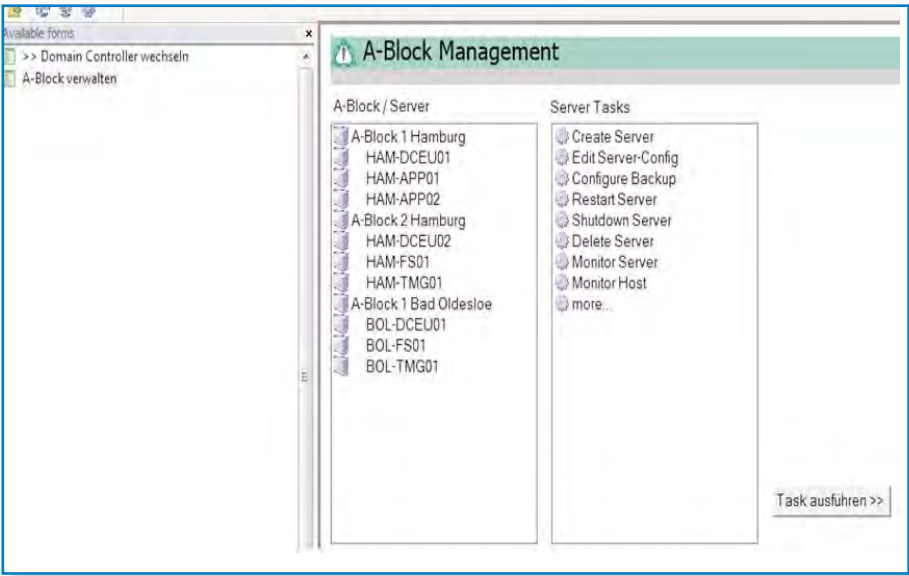


FIGURE 3: ASKLEPIOS IT-SERVICE MANAGEMENT TOOL FOR LOCAL IT TASKS

DEFINING AN EFFECTIVE CLOUD COMPUTING STORAGE STRATEGY

// Times change, as do our wills, what we are is ever changing; all the world is made of change, and forever attaining new qualities."

—Luíz Vaz de Camões

STORAGE REQUIREMENTS

The foundation of a cloud computing infrastructure is virtualization. Most physical challenges—such as underutilization of server resources, difficulty in protecting server availability, and disaster recovery—are all made easier with virtualization.

Because of the complexities associated with hypervisor management resources and the shared storage model, the biggest challenge in creating a cloud infrastructure is storage management.

In a cloud computing environment, there are usually two possible approaches to designing a storage solution: scale up and scale out.

Deciding which strategy to use will affect the overall cost, performance, availability, and scalability of the entire infrastructure. Defining the right direction based on your specific requirements is key for a successful cloud deployment.

Considering the pros and cons of each approach can help you make a better decision about which model to use, or whether to use a mix of the two models, to make your decisions easier.

CLOUD STORAGE

One objective of cloud computing is to be able to abstract the physical layer and manage the infrastructure based on policy and service definitions. However, to reach this objective, you must have an infrastructure that is well designed and prepared to scale based not only on the quantity, but also on the quality of compute components.

You can gain the benefits of infrastructure as a service (IaaS) with a large-scale deployment where infrastructure is shared among different kinds of workloads such as low latency and high throughput of the online transaction processing system. At the other extreme is



**THE BIGGEST
CHALLENGE
IN CREATING
A CLOUD
INFRASTRUCTURE
IS STORAGE
MANAGEMENT.**

STANDARDIZING IT

high-tolerance latency, backup, and archive deployment, where the amount of disk space is more important than speed (Figure 1).

To reach this level of flexibility in a seamless infrastructure, design it by standardizing as much as possible on the same fabric (i.e., storage area network [SAN] or network-attached storage [NAS]). It's essential to organize and automate storage tiering for an efficient cloud infrastructure. IT administrators must have predictive information on capacity and be able to make decisions about platform growth based on accurate information, allowing a granular investment in infrastructure without incurring service-level agreement (SLA) penalties.

ARCHITECTURE SELECTION

Defining the best infrastructure for a particular use case normally

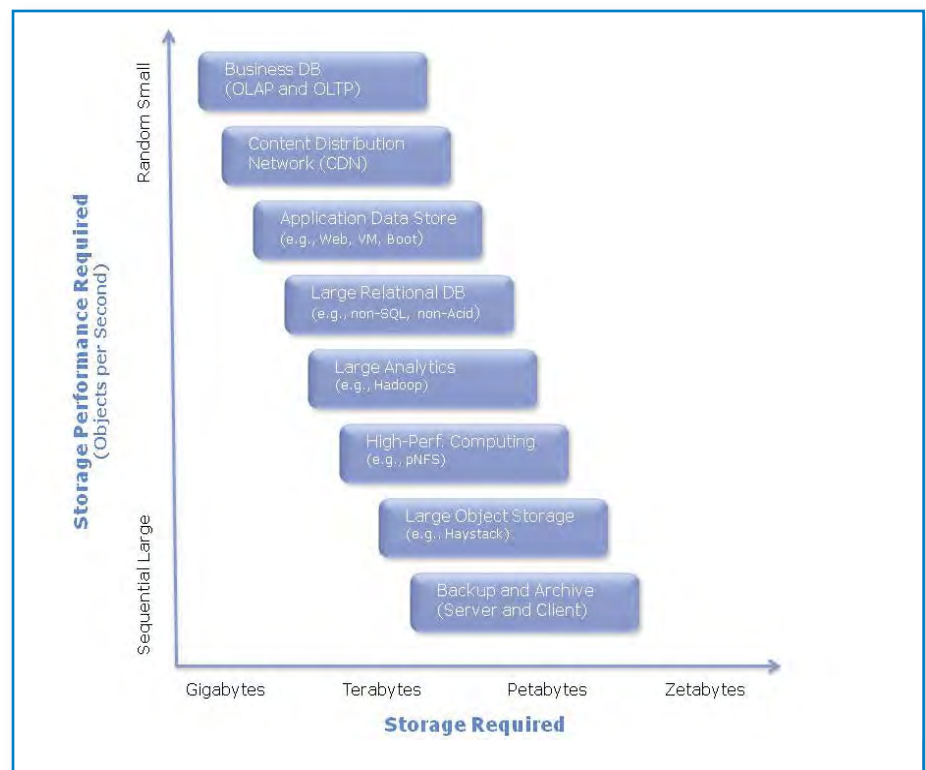


FIGURE 1. STORAGE REQUIREMENTS BASED ON WORKLOAD

means defining the technical requirements and finding the best solution in the marketplace that both meets those requirements and fits into the project's budget.

However, working with cloud computing, we usually deal with unknown numbers of variables.

Defining a common infrastructure for both actual and future services is an exercise in guessing.

The more we know about our environment, the better our decisions will be in the architecture design phase. Cloud computing can put us in an uncomfortable position, since we must make decisions without knowing which future applications and services it will need to support. There isn't a magic way to make the right choices. However, we can rely on information we do know and make the safest possible decisions.

DESIGNING STORAGE

To guide us through the process of architecture selection, Table 1 shows both the pros and cons of the scale-out and scale-up approaches.

Scale-out storage usually fits better in environments where you need to increase capacity with a low total cost of acquisition (TCA) and in small increments instead of making a major investment in scale-up storage. You can usually have the convenience of paying as you enable the resources.

From the SLA perspective, scale-up may impose a performance penalty as you grow, since you have a static amount of cache, memory, and CPU in the storage shared among a number of host machines. As you add machines to this pool, increasing competition for the same storage resources, storage access becomes slower.

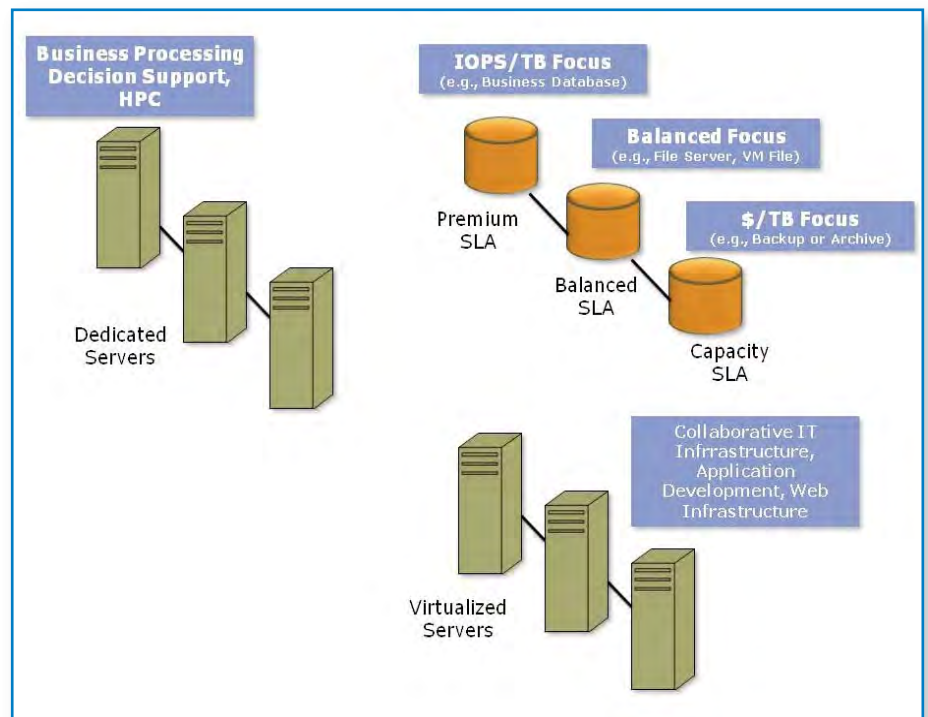


FIGURE 2. SAMPLE STORAGE ARCHITECTURE FOR CLOUD COMPUTING

With scale-out storage, as you add nodes, you also add CPU capacity, memory, cache, and spindles that are shared—potentially speeding access to data in storage.

From the management perspective, scale-up storage is much easier to maintain because of the centralized way it operates. Scale-out architec-

ture is harder to maintain since, depending on the scale-out solution you choose, you may mix nodes with different ages in the same storage cluster.

Each approach brings advantages and disadvantages. The role of the architect is to create a balance and find the best tradeoffs.

EFFECTIVE STORAGE STRATEGY

SOLUTION DESIGN

Independent of architecture selection, the reality for most organizations is dealing with a physical, dedicated server side-by-side with automated virtualized resources.

At the same time, you must provide enough storage space for backup and archive using disks with the best possible price/performance. From the other side, dedicated enterprise resource planning (ERP) systems and large online transaction processing

(OLTP) databases need transactions to take place as quickly as possible. A balanced storage strategy may be best for resources such as file servers and VM file storage (Figure 2).

Adopting a strategy of unifying networks (i.e., using the same fabric for storage and LAN access, with 10GbE interface), this design can be much easier to adapt and change as the organization adds new node and storage technologies.



TABLE 1. SCALE-OUT AND SCALE-UP COMPARISON

	Scale-Out (SAN/NAS)	Scale-up (DAS/SAN/NAS)
Hardware scaling	Add commodity devices	Add faster, larger devices
Hardware limits	Scale beyond device limits	Scale up to device limit
Availability, resiliency	Usually more	Usually less
Storage management complexity	More resources to manage, software required	Fewer resources to manage
Span multiple geographic locations	Yes	No

BALANCED APPROACH

Defining an effective storage strategy in a cloud infrastructure is key for a successful implementation. Mistakes can be very expensive to fix—and can ruin the quality

and price structure in a competitive marketplace such as IaaS.

For the best results as you define your storage strategy, it's essential

to consider not only a balanced scale-out and scale-up approach, but also your media access (fabric) strategy.

Learn more about cloud storage technologies and tomorrow's cloud [here](#).

[Back to Contents](#)

CLOUD SECURITY

Securing the Infrastructure with Intel® Trusted Execution Technology

Mark Wright

Senior Solution Architect, Intel Corporation

mark.a.wright@intel.com



Intel® Trusted Execution Technology (Intel® TXT) is a key security component for the cloud that provides hardware-based security technologies. It hardens platforms from the emerging threats of hypervisor attacks, BIOS, or other firmware attacks, malicious root kit installations, or other software-based attacks.

SECURING SENSITIVE DATA

It increases protection by allowing greater control of the launch stack through a measured launch environment (MLE) and enabling isolation in the boot process. It extends the virtual machine extensions (VMX) environment of Intel® Virtualization Technology (Intel® VT), permitting a verifiably secure installation, launch, and use of a hypervisor or operating system (OS).

As you evaluate Intel TXT for your own enterprise cloud environment, there are some key questions to ask:

- **WHO** needs Intel TXT?
- **WHAT** can it do?
- **WHAT** are the requirements for Intel TXT?
- **WHAT** are the use cases for Intel TXT?

Let's discuss each of these points and review some real-world lab installation and configuration pain points in setting up this type of security infrastructure to make it easier to provide cloud infrastructure security.

WHO NEEDS INTEL TXT?

Based on customer feedback, organizations that are seriously evaluating Intel TXT include financial institutions, pharmaceutical companies, and government agencies. These sectors must secure both sensitive data and the server hardware that supports it.

Security is the leading concern for IT groups implementing both private and public cloud solutions. To secure the cloud environment, and to have both the private and public sectors continue migrating to the cloud, it's essential to address security at all levels.

Intel believes it has developed a way to secure a key ingredient of the hardware platform, system boot of the hardware infrastructure.

One key question around security is how can we know when we launch our systems—especially in the



**SECURITY IS
THE LEADING
CONCERN FOR
IT GROUPS
IMPLEMENTING
BOTH PRIVATE
AND PUBLIC
CLOUD
SOLUTIONS.**

TRUSTED BOOT

cloud—if they are secure? Wouldn't it be nice to verify the programs and data you entrust to a cloud provider are running on securely-booted hardware? Knowing that no malware or rootkit has injected itself into your cloud infrastructure is very valuable.

Once you've provided a secured boot, the next step is to provide runtime integrity checking of your cloud infrastructure, ensuring your applications running in the cloud continue to be verified as secure. Currently, Intel TXT provides security functions but not ongoing security protection. This will be the next step in securing the cloud.

WHAT CAN INTEL TXT DO?

Intel TXT provides a trusted boot (tboot), the ability for a virtual environment or virtual machine manager (VMM) module to validate that when it boots, it is secure, using dynamic root of trust measurement (DRTM). Intel TXT provides this capability through an infrastructure based in

the Intel® Xeon® processor and known as the root of trust. Intel TXT checks the consistency in behaviors and launch-time configurations against a verified benchmark called a known good sequence. The system can then quickly assess and alert against any attempts to alter or tamper with a system's launch-time environment.

WHAT ARE THE REQUIREMENTS FOR INTEL TXT?

Intel TXT includes these components:

- **INTEL XEON PROCESSOR.** Intel Xeon processors 5600 series and beyond include Intel TXT and Intel VT-capable silicon for the root of trust
- **INTEL CHIPSET WITH INTEL VT** that provides the isolation capabilities for measured launch
- BIOS. Intel TXT support from within BIOS
- **AC MODULE** created and signed by Intel inside the BIOS

- **TRUSTED PLATFORM MODULE (TPM)** integrated onto the motherboard that provides securely-generated cryptographic keys
- **SINIT**, the instruction set for initiating a secure launch of VMM or the OS
- **VMM**, an Intel TXT-aware hypervisor

There are many moving parts that make Intel TXT work within a server. But there doesn't seem to be a good, publicly-available utility to validate that all these components are available and functional for server systems. The only way to find out is to verify with your OEM that a particular system is Intel TXT-compliant and contains all required components.

INSTALLING THE COMPONENTS

To set up a sample Intel TXT-capable system in a lab environment, we used:

- **A NEW SERVER** with an Intel Xeon processor 5600 series and Intel chipset that supported Intel TXT

SETTING IT UP

- **ENOUGH MEMORY** and hard disk drive space to support our virtual environment
- **A HYPERVISOR**, in this case VMware* 5.1
- **A HYTRUST* SECURITY APPLIANCE** that installs as a virtual machine in our virtual environment to validate and enforce our secure environment

To learn more, check this [compatibility table](#) that shows supported OEMs, ISVs, operating systems, and hypervisors that are Intel TXT-aware.

BIOS SETUP

First go into the BIOS, under the processor configuration, and select Intel TXT. Enable and set the admin password.

Also in BIOS, under security, enable TPM to “on” and “functioning.” Note that it looks like nothing happened in BIOS; it just takes you to the previous screen. Go back to verify that TPM is set. Save the settings and reboot the system.

HYPERVISOR SETUP

Install and set up VMware 5.1 and your virtual machines (VMs) that will be used for VMotion operation.

Make sure you have the latest HyTrust appliance 2.5.3. The latest versions of VMware and HyTrust resolve a number of bugs in supporting Intel TXT.

Also, the dynamic resource scheduling (DRS) and dynamic power management (DPM) automation features within VMware will require secure VMotion policies to be set up, or VMotion disabled altogether, to make sure that VMs don’t migrate to untrusted hosts without going through the HyTrust appliance first. VMware was chosen because its hypervisor supports Intel TXT automatically. The trust is either automatically set up between Intel TXT or validated (signed) for a trusted launch with vSphere during the installation process. Not all hypervisors or operating systems behave in the same way.

As an example, if you were to create a TXT trust environment manually in Linux* you would have to go through these steps:

- **RUN** the tboot Installation package
- **DOWNLOAD** the SINIT ACM from Intel’s website
- **MOVE** the SINIT file to the /boot/ directory (\$ mv <SINIT-FILE> /boot/)
- **VERIFY** that you have all the latest files
- **RUN** the TCSD Daemon
- **INSTALL** the TCG software stack
- **MODIFY** the “GRUB” file to boot to the new tboot kernel
- **REBOOT** your system
- **VERIFY** that the platform configuration registers (PCRs) are populating and Intel TXT measured launch equals “true”

SECURITY APPLIANCE SETUP

We used the HyTrust appliance and added it as a VM within our VMware

MAXIMIZING VALUE

environment. The HyTrust configuration of our virtual environment and security policy configurations, which can be complex, took some time to set up.

WHAT ARE THE USE CASES FOR INTEL TXT?

Once you have created a secure server environment, it's important to consider all the possible uses of Intel TXT to maximize its value in your cloud environment:

- **VERIFIED LAUNCH.** Intel TXT allows you to ensure, upon system boot, that your system has launched into a trusted state.
- **TRUSTED POOLS.** Intel TXT allows you to create pools of vir-

tual resources that are hosted on trusted resources. You can then manage live migration of VMs between hosts, based on policies of trust, and constrain critical apps to run only on trusted platforms. You can also prevent selected apps from running on trusted platforms.

- **COMPLIANCE.** You can ensure reporting and audit compliance for mandated security control requirements for sensitive data.
- **ATTESTATION.** Geo tagging can let you constrain apps to run in selected countries, regions, or states as authenticated via PCRs.



SECURING YOUR DATA

SECURITY IN THE CLOUD: WHAT DID WE LEARN?

We learned that Intel TXT does enable a key component of security in the cloud environment. The early adopters have limited hardware and software support, but that will improve over time.

We found out that very few vendors are automated in setting up a secure boot environment. You need secure platform component validation tools

to verify TXT-capable systems and automation tools to improve supporting secure server provisioning.

The initial list of use cases for Intel TXT is compelling, and will continue to expand with the evolution from secure boot. As the transformation occurs from this first security step to an evolution of integrity-checking capabilities, secure infrastructure will begin to provide the continuous security needed to build a truly

trusted environment in the private and public cloud.

An evaluation of Intel TXT is an important first step toward building a trusted environment. This is important for financial institutions, pharmaceutical companies, and government entities, as well as for any company with sensitive data in the private or public cloud. Using tools that can help ensure your data is secure will enable the cloud to flourish in the business world.

To learn more about Intel TXT, visit www.intel.com/go/txt.

[Back to Contents](#)

Intel® Hyperthreading Technology requires an Intel® HT Technology enabled system, check with your PC manufacturer. Performance will vary depending on the specific hardware and software used. Not available on Intel® Core™ i5-750. For more information including details on which processors support HT Technology, visit <http://www.intel.com/info/hyperthreading>.

No system can provide absolute security under all conditions. Intel® Anti-Theft Technology requires an enabled chipset, BIOS, firmware and software and a subscription with a capable Service Provider. Consult your system manufacturer and Service Provider for availability and functionality. Intel assumes no liability for lost or stolen data and/or systems or any other damages resulting thereof. For more information, visit <http://www.intel.com/go/anti-theft>.

No computer system can provide absolute security under all conditions. Intel® Trusted Execution Technology (Intel® TXT) requires a computer system with Intel® Virtualization Technology, an Intel TXT-enabled processor, chipset, BIOS, Authenticated Code Modules and an Intel TXT-compatible measured launched environment (MLE). Intel TXT also requires the system to contain a TPM v1.s. For more information, visit <http://www.intel.com/technology/security>.

Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM). Functionality, performance or other benefits will vary depending on hardware and software configurations. Software applications may not be compatible with all operating systems. Consult your PC manufacturer. For more information, visit <http://www.intel.com/go/virtualization>.

Intel® Turbo Boost Technology capability requires a system with Intel Turbo Boost Technology capability. Consult your PC manufacturer. Performance varies depending on hardware, software and system configuration. For more information, visit <http://www.intel.com/technology/turboboost>

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at www.intel.com.

Copyright © 2011 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others. Printed in USA SS/PP/0712 Please Recycle

