

An Oracle White Paper
October 2011

Database Consolidation onto Private Clouds

Executive Overview	3
Cloud Overview and Considerations.....	4
Business Drivers for Consolidating Databases onto a Private Cloud .	6
Steps to Database Consolidation onto a Private Cloud.....	7
Models for Database Consolidation onto a Private Cloud	9
Choosing the Appropriate Cloud Deployment Model	12
Building Private Cloud Infrastructures.....	13
Oracle Technologies for Consolidating Databases on Private Clouds	14
Conclusion	17

Executive Overview

Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources or shared services (e.g., networks, servers, storage, applications, and IT services). The key benefits of cloud computing are reduced costs, reduced complexity, improved quality of service, and increased flexibility when responding to changes in workload.

In traditional models, business applications are bound to a particular infrastructure, resulting in low efficiency, utilization, and flexibility. Cloud computing technologies allow applications to be dynamically deployed onto the most suitable infrastructure at runtime. The elastic nature of cloud computing allows applications to scale at lower cost without requiring “fork-lift” upgrades.

IT departments, which are under increasing pressure to provide computing services at the lowest possible cost, can choose either public or private clouds to meet these needs. However, driven by concerns over security, regulatory compliance, control over quality of service, and long-term costs, many customers choose internal private clouds.

Private clouds provide the same cost and flexibility benefits as public clouds, but they also enable IT departments to control the quality of service delivered to their users. In addition, private clouds allow IT departments to better secure data and meet governance regulations—a major concern when storing data in external, public clouds.

This paper describes how Oracle Database11g can be successfully consolidated onto a private cloud through several deployment models.

Cloud Overview and Considerations

According to the U.S. National Institute of Standards and Technology, cloud computing consists of five essential characteristics, three distinct service models, and four deployment models.

Essential Characteristics	Service Models	Deployment Models
On-Demand Self-Service	Software-as-a-Service (SaaS)	Public Cloud
Resource Pooling	Platform-as-a-Service (PaaS)	Private Cloud
Rapid Elasticity	Infrastructure-as-a-Service (IaaS)	Community Cloud
Measured Service		Hybrid Cloud
Broad Network Access		

Service Models

Software-as-a-Service

This term generally refers to applications that are delivered to end users over the Internet or broad band access. There are hundreds of SaaS providers covering a wide variety of applications. Oracle CRM On Demand, Salesforce.com, and Google Apps are examples of the SaaS model.

Infrastructure-as-a-Service

Refers to computing hardware (servers, storage, and network) delivered as a service. This typically includes the associated software as well, including operating systems, virtualization, clustering, and so on. Amazon Web Services, for example, offers their Elastic Compute Cloud (EC2) for compute servers, SimpleDB for database, and Simple Storage Service (S3) for storage.

Platform-as-a-Service

This model conveys how an application development and deployment platform can be delivered as a service to developers, allowing them to quickly build and deploy an SaaS application for end users. These platforms typically include database and middleware, and are often specific to a language or API. For example, Google AppEngine is based on Java and Python, EngineYard is based on Ruby on Rails, and Force.com uses a proprietary variation of Java.

Database Services on a Private Cloud

For database environments, the PaaS cloud model provides better IT services than the IaaS model. The PaaS model provides enough resources in the cloud that databases can quickly get up and running and still have enough latitude for users to create the applications they need.

Additionally, central IT management, security, and efficiency are greatly enhanced through consistency and economies of scale. Conversely, with the IaaS model, each tenant must build most of the stack on their own, lengthening time to deployment and resulting in inconsistent stacks that are harder to manage.

A private cloud is an efficient way to deliver database services because it enables IT departments to consolidate servers, storage, and database workloads onto a shared hardware and software infrastructure. Databases deployed on a private cloud offer compelling advantages in cost, quality of service, and agility by providing on-demand access to database services in a self-service, elastically scalable, and metered manner.

Private clouds are a better option than public clouds for many reasons. Public clouds typically provide little or no availability or performance service-level agreements, and there are potential data security risks. In contrast, private clouds enable IT departments to have complete control over the performance and availability service levels they provide, and can easily enforce data governance regulations and auditing policies.

Business Drivers for Consolidating Databases onto a Private Cloud

Four key business drivers typically motivate database consolidation onto a private cloud.

Reduced Cost

IT budgets are under constant scrutiny, so IT departments need solutions that reduce both capital expense and operating expense without compromising key business requirements. Consolidating shared resources effectively replaces siloed, underutilized infrastructures with a shared resource pool, which lowers overall costs and increases resource utilization. Capital expenditure can be reduced beyond simply shrinking server footprint by creating a higher density of databases per server through multi-tenancy configurations. And operational expenditure can be reduced by improving efficiency through automation and improved management productivity.

Reduced Complexity

IT departments can simplify their environments by reducing the number of supported configurations and services through rationalization, standardization, and consolidation. By standardizing on a common set of building blocks, IT departments can easily deploy predefined configurations and scale-out using modular components. One of the keys to reducing complexity is centralized management: as the environment becomes more homogenous, it becomes easier to manage. And having a central management hub keeps operational costs low and further promotes the automation of routine tasks.

Increased Quality of Service

IT departments are not only trying to drive down costs, they are also looking for solutions that will improve performance, availability, and security. In a private cloud, database performance can be monitored and managed via shared Centers of Excellence. Databases also benefit from the high availability built into the private cloud. And consolidation helps enforce a unified identity and security infrastructure as part of standardized provisioning process.

Improved Agility

IT departments are increasingly looking to develop more agile and flexible environments that will enable faster time to market and a rapid response to changing business requirements. This will provide efficient rollout of new business strategies as well as the capability to quickly deploy applications without a huge lag time due to infrastructure setup. The three key aspects of agility are

- **Fast deployment.** Building a private cloud infrastructure using standard hardware components, software configurations, and tools enables an automated and simplified deployment process.

- **Rapid provisioning.** Resources in a cloud can be rapidly provisioned, often via self-service, providing quicker application deployment. This reduces overall time in deploying production applications, development platforms, and creating test bed configurations.
- **Resource elasticity.** The ability to grow and shrink the capacity of any database, both in terms of size and compute power, offers applications the flexibility to meet the dynamic nature of business workloads.

Steps to Database Consolidation onto a Private Cloud

Building a private cloud requires the transformation and optimization of the IT infrastructure, and that is typically executed in three steps: rationalization, architecture optimization, and implementation of shared services.

Rationalization

IT rationalization determines the best use of IT services and reduces nonproductive redundancy throughout the enterprise. IT departments should rationalize their technology architecture by standardizing their service portfolio and technology stack. Through standardization, the IT environment becomes much more homogenous, which makes it easier to manage. It also reduces costs and complexity and increases agility.

Architecture Optimization

All layers of the technology stack must support service-level objectives and growth requirements. Scalability, availability, data security, and datacenter management are only as strong as the weakest link. Balanced technology architectures employ virtualization, consolidation, and management automation to meet business requirements. Virtualization, for example, transforms the typical server-to-application silo model to a multi-tenancy model. The key to virtualization is not necessarily the underlying technology, but rather the capability to abstract resources requested by the business from resources fulfilled by IT.

Shared Services

IT departments can leverage shared services to reduce costs and meet the demands of their business users, but there are many operational, security, organizational, and financial aspects of shared services that must be managed to ensure effective adoption.

Consolidation is vital to shared services, as it allows IT to restructure resources by combining multiple applications into a cohesive environment. Consolidation goes beyond hard cost savings; it simplifies management, improves resource utilization, and streamlines conformity to security and compliance standards. Therefore, the next item to consider is the level of consolidation that can be achieved in a private cloud architecture.

- **Server Consolidation.** Reduce the number of physical servers and consolidate databases onto a smaller server footprint.

- **Storage Consolidation.** Unify the storage pool through improved use of free space in a virtual storage pool.
- **Database Consolidation.** Reduce the number of operating system installations. Reducing server footprint does not always provide the best ROI, but reducing the number of operating systems will improve overall manageability.
- **Schema Consolidation.** Reduce the number of database instances through schema consolidation. Separate databases are consolidated as schemas in a single database, reducing the number of databases to manage and maintain.
- **Extreme Schema Consolidation.** Merge the redundant databases that support business intelligence or operational data store systems. By consolidating into a single data store, these workloads benefit from the additional resources and scalability provided by the private cloud infrastructure.

As shown in Figure 1, the potential return on investment (ROI) increases as the level of consolidation onto a private cloud increases. Cost savings, management efficiency, and improved resource utilization increase as server, storage, databases, schemas and workloads are consolidated.

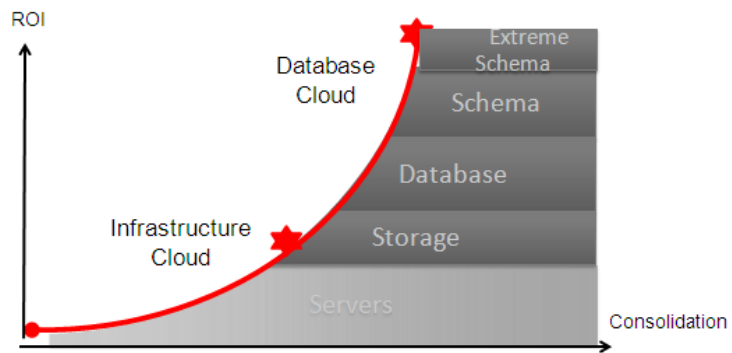


Figure 1. Consolidation relation to return on investment

Models for Database Consolidation onto a Private Cloud

Consolidating databases onto a private cloud is typically done in one of two ways: Infrastructure Cloud or Database Cloud.

Infrastructure Cloud

This deployment model is enabled via virtualization: servers are added to a server pool, and virtual machine (VM) guests are created from servers in the pool. There is generally a one-to-many relationship between servers and VM guests, with density driven by physical server resources (CPU and memory). When a database service is requested, the entire operating system stack is built and provisioned as part of VM guest creation.

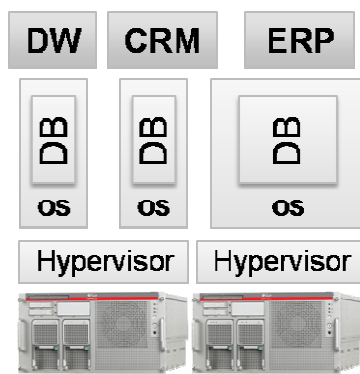


Figure 2. The infrastructure cloud/server consolidation model

In this model (illustrated in Figure 2), tenant isolation is at the VM guest level. Databases deployed will generally be single-instance databases that are instantiated in the VM guest. Scalability in this model is vertical, in that it is provided by adding more virtual resources (CPU or memory). VM guests, however, cannot span across servers in the server pool, meaning that the full resources of the private cloud cannot be brought to bear on a workload requirement. They can leverage the high availability capabilities provided by virtualization technology such as Live Migration or HA Restart.

The primary drivers for adopting this model include

- Straightforward, “push-button” database deployment via VM templates or profiles
- Support for provisioning of multiple database versions and configurations
- Excellent fault and resource isolation at the VM guest level
- Most applications run seamlessly against databases deployed in VM guests

Deployment considerations include

- Virtualization doesn’t reduce the number of operating systems or databases, therefore there is less actual consolidation and a lower return on investment
- This model encourages operating system and database sprawl, which creates a more complex environment to manage as the number of deployments increases

- This model is not as high-performing as other deployment models, and I/O-intensive databases may not perform well in virtualized environments
- Scalability is limited to the size of the biggest server in the pool, and database workloads cannot be scaled horizontally across the pool
- Resource management at the VM guest layer provides coarse-grain resource and workload management on a single machine, which limits ability to meet changing workload requirements

Database Cloud

A database cloud refers to database deployments in virtual pools using server clustering. There are two ways to deploy: provision databases natively or provision a schema to a shared database.

Provision Databases Natively – Database Consolidation

By standardizing on a common operating system, any one server from a cluster can host one or more database instances accessible via named database services. In this deployment model, a database is provisioned onto physical servers that are clustered together in a private cloud. By utilizing Oracle Real Application Clusters (Oracle RAC) or Oracle RAC One Node, databases inherit high availability through server redundancy. Elasticity and scalability is provided by adding additional nodes to the server pool (scaling out) or by adding more physical resources such as CPU, memory, or I/O cards to an existing node (scaling up).

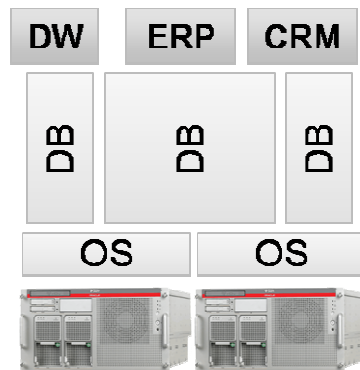


Figure 3. The Database Cloud- Database consolidation model

In this model (illustrated in Figure 3), tenant fault isolation is provided at the node or Oracle RAC service level. Oracle Database 11g features such as Instance Caging, Database Resource Manager, Automatic Workload Management, and Quality of Service provide tenant performance isolation. This enables efficient use of shared application resources to meet service level requirements.

The primary drivers for adopting the operating system consolidation model include

- High database to server consolidation density without incurring operating system sprawl
- Reduced server footprint with increased storage and server utilization, providing a higher return on investment
- Databases are not inhibited by virtualization layers, so they can deliver higher performance
- All applications are supported, and fine-grained workload management can be enforced

Deployment considerations include

- Standardization on a common operating system is required

Provision a Schema to a Shared Database – Schema Consolidation

In this deployment model (illustrated in Figure 4), the consolidated database essentially consists of one or more application schemas running across one or more servers in a private cloud, using Oracle Database 11g and Oracle Real Application Clusters.

In this model, when a tenant requests a database to be provisioned, a new schema is created, with its own set of tablespaces and corresponding service name. Oracle Real Application Clusters provides high availability and scalability across the private cloud. In this example, DW, ERP, and CRM are different schemas within a consolidated database, each with its own named database service spanning across server resources in the private cloud. Resource isolation is provided at the service level, and the Automatic Workload Management and Quality of Service features ensure that sufficient cloud resources are available to meet application performance and high availability requirements, even as the workload changes.




Figure 4. The Database Cloud - Schema consolidation model

Primary drivers for adopting of the database consolidation model include

- Provisioning a database service via a schema is extremely fast
- Consolidation into a larger database reduces operating system and memory overhead
- By reducing the number of server, operating system, and database deployments, this model enables dense consolidation provides the highest return on investment

Deployment considerations include

- Standardization on a common operating system and database version is required
- Some applications have stringent schema naming standards, and schema consolidation is vulnerable to namespace collisions, therefore application certification may be required
- There is limited resource, fault, and security isolation; however, technologies such as Oracle Database Vault and Oracle Database 11g's Quality of Service Management feature can minimize these isolation impacts

Choosing the Appropriate Cloud Deployment Model

In a private cloud, a database appears as a database service that users can easily access to read or write data and run their business applications. Users need not be concerned with management of the consolidated private cloud infrastructure, but they will require stringent performance, availability, and data security service levels. IT departments must choose the best deployment model to meet both budgetary constraints and business application service levels. Consolidation onto a private cloud using Oracle Database 11g offers a choice of deployment models. The table below summarizes the capabilities and benefits of each model.

	SERVER CONSOLIDATION	DATABASE CONSOLIDATION	SCHEMA CONSOLIDATION
IMPLEMENTATION	Easy	Easy	Difficult
APPLICATION SUITABILITY	Some	All	Some
ISOLATION	Highest	High	Limited
AVAILABILITY	High	Highest	Highest
SCALABILITY	Limited	Excellent	Excellent
CONSOLIDATION DENSITY	Low	High	Highest
RETURN ON INVESTMENT	Low	High	Highest

Although the server consolidation model provides ease of migration and high levels of isolation, it has performance limitations, lower availability, and the lowest return on investment. Database cloud models provide higher levels of consolidation than an infrastructure cloud. More importantly, an infrastructure cloud only consolidates servers and storage, while a database cloud consolidates servers, storage, operating systems, database versions, and workloads, offering a much higher return on investment.

Building Private Cloud Infrastructures

Once you've selected the appropriate consolidation model, the next item to consider is building the physical private cloud infrastructure. IT departments effectively have three choices.

Build Your Own

This is the traditional approach—the IT department puts everything together, including the rationalization, standardization, and consolidation phases; the definition of the server, storage, and network architecture; and the selection of the deployment model.

Reference Configurations

This configuration consists of a predefined and preconfigured server, storage, and network architectures. Validated reference configurations can reduce the onus on IT departments to define and configure systems. In fact, reference configurations can be superior to one you build on your own, but typically they still require significant installation and configuration.

Oracle Engineered Solutions

Oracle offers several Engineered Solutions, which combine hardware and software into complete, integrated and fully tested products that address the full range of database deployments for enterprise computing. Oracle Exadata Database Machine, Oracle Database Appliance, and Oracle SPARC SuperCluster are solutions designed to be optimal platforms for Oracle Database, and therefore an ideal platform for Private Database Cloud computing.

Oracle Exadata Database Machine

Oracle Exadata Database Machine implements all the hardware and software required for native deployments of Oracle Database 11g onto a private cloud. It combines servers, storage, and networks into one integrated package, eliminating difficult integration problems. It also provides highly optimized performance capabilities not available through the other infrastructures. Rather than going through the entire rationalization and standardization process, IT departments can simply implement Oracle Exadata Database Machine for database consolidation onto a private cloud.



Figure 5. Oracle Exadata Database Machine X2-8

Oracle Database Appliance

The Oracle Database Appliance is a complete plug-n-go system that offers customers a fully integrated system of software, servers, storage and networking in a single box, delivering high availability database services for a wide range of home grown and packaged OLTP and Data Warehousing applications. Oracle Database Appliance is pre-installed with Oracle Linux and Oracle Appliance Manager software, saving time and money by simplifying the deployment, maintenance, and support of highly available database solutions. The Oracle Database Appliance is the perfect clustered database solution for rapid deployment and maximum uptime.




Figure 6. Oracle Database Appliance

Oracle SPARC SuperCluster

The Oracle SPARC SuperCluster delivers high performance, availability, scalability and security to a wide range of enterprise applications, including database, middleware and Oracle and custom applications. An optimized package of servers, storage and software, SPARC SuperCluster integrates Oracle Exadata and Oracle Exalogic Elastic Cloud with SPARC T4-4 servers, ZFS Storage Appliance, InfiniBand I/O fabric, and Oracle Solaris. With the support for both Solaris 11 and Solaris 10, the opportunity to offer cloud services for new and legacy database and application versions is a compelling differentiator.



Figure 7. Oracle SPARC SuperCluster

Oracle Technologies for Consolidating Databases on Private Clouds

Over the last decade and more, Oracle has introduced a variety of innovations that help IT departments consolidate databases onto private clouds. For example, Oracle Real Application Clusters enables server resources to be shared on a cluster, Oracle Automatic Storage Management enables storage resources to be virtualized and shared, and the Database Resource Manager and Instance Caging features in Oracle Database automate resource optimization—just to name a few. The innovation continues with Oracle Database 11g server pooling, the Quality of Service Management feature.

Oracle Real Application Clusters

Oracle Real Application Clusters supports the transparent deployment of databases across pools of servers, providing fault tolerance from hardware failures or planned outages. Support is provided for custom-developed online transaction processing (OLTP) and data warehouse applications, as well as popular packaged products such as SAP, Oracle's PeopleSoft, Oracle's Siebel, and Oracle E-Business Suite applications. Oracle RAC provides high availability for databases by removing the single server as a single point of failure. If a node in a server pool fails, Oracle Database continues to run on the remaining servers in the pool. Oracle Real Application Clusters provides great flexibility for scaling databases across private cloud infrastructures. In order to keep costs low, private clouds can be built on standardized, commodity-priced servers, storage pools, and network components. When more processing power is needed, another server can be added without taking users offline.

Oracle Automatic Storage Management

Oracle Automatic Storage Management provides file system and volume management functionality for Oracle Database files. Oracle Automatic Storage Management is easier to manage than conventional file systems, has the performance of raw volumes, and is tightly integrated with Oracle Database 11g, forming the foundation of storage virtualization. In a private cloud infrastructure, Oracle Automatic Storage Management simplifies storage consolidation and provisioning, and provides an easy way to migrate existing databases to a private cloud infrastructure. Oracle Automatic Storage Management also includes an innovative rebalancing capability that distributes data evenly across all storage resources in the private cloud, providing an even distribution of I/O for all databases consolidated on the storage pool. Oracle Automatic Storage Management also provides high database availability with automatic mirror reconstruction and resynchronization.

Diagnostic and Tuning Packs

Oracle Diagnostics Pack offers a comprehensive set of automatic performance diagnostics and monitoring functionality built into the core database engine. Oracle Diagnostics Pack offers a complete, cost-effective, and easy-to-use solution for managing the performance of Oracle Databases deployed onto a private cloud infrastructure. When used with Oracle Enterprise Manager, Oracle Diagnostics Pack also provides enterprise wide performance and availability reporting, a centralized performance repository, and valuable cross-system performance aggregation, significantly simplifying the task of managing large sets of databases across a private cloud.

Oracle Tuning Pack is an extremely cost-effective and easy-to-use solution that automates the entire application tuning process. Enhancement of SQL performance is achieved through real-time monitoring and SQL advisors that are seamlessly integrated with Oracle Enterprise Manager. Together they provide a comprehensive solution for automating the complex and time-consuming task of application tuning.

Workload and Quality of Service Management

Oracle Database 11g provides capabilities that allow workloads to be isolated to the databases that are consolidated on a private cloud. Individual databases can be caged to run on specified cores within a private cloud, allowing for dense consolidation. Within a database, different application users can be mapped to database wide resource plans that govern how much memory, CPU, and I/O bandwidth (on Oracle Exadata Database Machine) each user consumes, making it easy to constrain resource usage in a shared environment.

Database Provisioning and Configuration Management

Oracle Provisioning and Patch Automation Pack automates the deployment of new databases to the private cloud. Gold images can be defined for different types of database deployments, or existing database environments can be cloned and then provisioned to a private database cloud from a central management console. Once these standard configurations have been deployed, the solution also identifies and downloads any required patches to keep the configurations up-to-date with fixes to any reported issues.

Oracle Configuration Management Pack provides a simplified way to view the configurations of all databases and related components across a private cloud. It can centrally define baselines for these configurations, ensuring that the deployments do not deviate from the defined configurations.

Conclusion

Consolidating databases onto a private cloud is a new model for the delivery of database services. Private clouds consolidate servers, storage, operating systems, databases, and mixed workloads onto a shared hardware and software infrastructure. Deploying databases on a consolidated private cloud enables IT departments to improve quality of service levels—as measured in terms of database performance, availability, and data security—and reduce capital and operating costs. Consolidation can be achieved through server, operating system, and database consolidation, and the higher the consolidation density achieved, the greater the return on investment.

Oracle Database 11g and Oracle Real Application Clusters enable all levels of consolidation, offering IT departments a choice of private cloud deployment architectures. Oracle Database 11g and Oracle Real Application Clusters are key software components of Oracle Exadata Database Machine, Oracle Database Appliance and Oracle SPARC SuperCluster, making each an excellent private cloud database consolidation platform. These solutions provide pre-integrated configurations of hardware and software components engineered to work together, and optimized for different types of database workloads. Choosing the best platform for your organization will be one of the key milestones in your roadmap to realizing the benefits of Cloud Computing for your private databases.



White Paper Title
February 2011 – v45
Author: Nitin Vengurlekar
Contributing Authors: Michael Timpanaro-
Perrotta, Mark Macdonald, Burt Clouse, Philip
Newlan

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 1010

Hardware and Software, Engineered to Work Together