IT@Intel White Paper

Intel IT
IT Best Practices
Cloud Computing and Compute Continuum
June 2012

Client-aware Services in the Cloud

Building on our previous exploration of client-aware cloud technologies and the development of several compelling usages associated with distributing workloads across the client and cloud, we are now implementing foundational client-aware services that support our compute continuum and enterprise private cloud efforts.

Executive Overview

Intel IT is transforming our application delivery model to enable the back-end cloud and the front-end client to work together to support an increasing number of combinations of operating systems, devices, and computing models. Our goal is for users to have seamless, secure access to cloud-based corporate and personal applications and services across the broad range of devices that make up the compute continuum. Applications that can recognize and make use of client capabilities can make this vision a reality.

Building on our previous exploration of client-aware cloud technologies and the development of several compelling usages associated with distributing workloads across the client and cloud, we are now implementing the following foundational client-aware services that support our compute continuum and enterprise private cloud efforts:

- **Security service.** Handles separate levels of access for different devices based on malware prevention, device management, and content protection.
- **Content synchronization service.** Enables computing from anywhere by synchronizing content settings and states between devices and the cloud. Also supports end-user collaboration by enabling content sharing.
- **Business application service.** Creates and uses robust, secure, context-aware, and device-aware applications across platforms with access to enterprise legacy back-end services.

These foundational services tie the client and the cloud together, taking advantage of the strengths of each.

Omer Ben-Shalom
Senior Principal Engineer, Intel IT

Chuck Brown
Director, Emerging Compute Lab
Intel Architecture Group

John Dunlop
Enterprise Architect, Intel IT

Contents

IT@INTEL

The IT@Intel program connects IT professionals around the world with their peers inside our organization – sharing lessons learned, methods and strategies. Our goal is simple: Share Intel IT best practices that create business value and make IT a competitive advantage. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BACKGROUND

The biggest challenge we have experienced with developing a client-aware enterprise private cloud is addressing the diversity and complexity that comes with an increasing number of operating systems, devices, and computing models. Intel IT's enterprise private cloud environment is deployed across data centers worldwide; 80 percent of new business services are deployed within the cloud. And through the transformation of our enterprise security architecture, we enable a wide range of devices and application delivery models while having to protect Intel's critical assets.

As we move toward a cloud-centric application delivery model, we are focused on providing the best possible user experience. Applications need to recognize and make use of client capabilities—thereby improving user experience—while enabling IT to maintain the information security and manageability necessary to meet legal obligations and to protect users' privacy. To this end, we have been investigating client-aware technologies, such as mobile application frameworks, rich Internet applications (RIAs), and workload shifting. These types of technologies help enable central workspace management in the cloud as well as the best possible user experience across devices.

We have also identified several compelling usages associated with distributing workloads between the client and the cloud.

- **Instant collaboration.** The ability to easily initiate audio, video, and data conferencing, with one or more team members from any notebook, smartphone, or tablet and collaborate on content.
- **Adaptive workspace.** The use of multiple personalized profiles to dynamically adapt to the environment. The settings, appearances, and services match the employee's needs whether at work, home, or traveling.
- **Business assistant.** The ability to provide different types of assistance to employees by taking advantage of the user and device context. One example is enabling an employee who is visiting an unfamiliar Intel campus to quickly find available local corporate resources such as a conference room, printer, or work station.

With an understanding of several client-aware technologies and potential enterprise usages, our next step is to implement client-aware services that support a wide variety of devices accessing cloud-based applications.

SOLUTION

We are working to implement foundational capabilities that enable the back-end cloud and the front-end client to work together. These capabilities will form a platform-as-a service that includes elements of both the device and the cloud.

We have identified three client-aware services, shown in Figure 1, that are

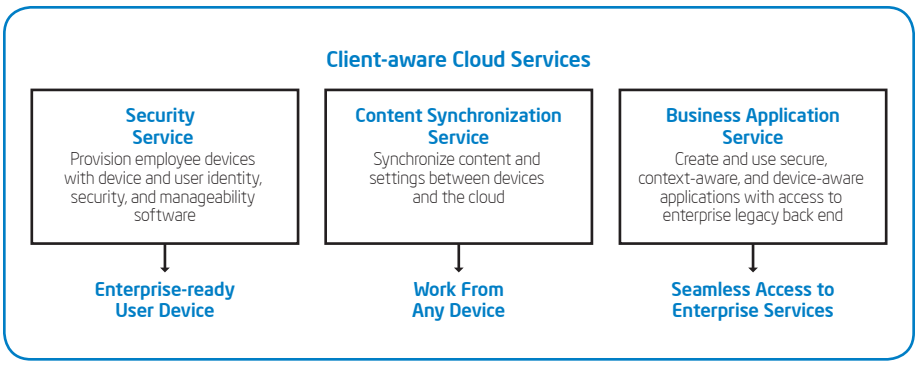

Figure 1. Security, content synchronization, and business application services form the foundation for a client-aware cloud.

critical to implementing a client-aware cloud.

- **Security service.** Establishes trust that is device-aware and provides device management and content protection.
- **Content synchronization service.** Enables computing from anywhere by synchronizing content, settings, and states among devices and the cloud. It also supports content sharing for collaboration.
- **Business application service.** Creates and uses robust, secure, context-aware, and device-aware applications across platforms with access to enterprise legacy back-end services.

Without these services, fully supporting IT consumerization and providing seamless access to services and applications across a wide variety of devices—that is, across the compute continuum—will be challenging for enterprises and provide a poor experience for end users. These foundational services tie the client and the cloud together, taking advantage of the strengths of each and making it easier to create software that runs across the variations of form factors and operating systems.

Having defined the foundational capabilities required for establishing a client-aware cloud, we are beginning to develop these services. Because these services are not yet available in the market, we have chosen to initially blend third-party and custom solutions.

Security Service

Although we are implementing centrally managed cloud services, we need to also be aware of the client, so we can protect data at rest, in use, and in transit, and differentiate between different devices that have different feature sets. As shown in Figure 2, the security service addresses the level of trust, device manageability, and content protection.

DETERMINING LEVEL OF TRUST

In our environment, we have found that it simply isn't practical to deliver the same set of services to every device. There is a large variety of different hardware, OSs, and ownership models, and different levels of controls make some devices more capable than others. Some devices do not have the features necessary to meet IT's minimum security configuration for certain levels of confidential data; only a subset of devices can gain wide access to corporate data. We build our decisions for access control and authorization in part on the trustworthiness of the device. In addition, we consider how the diversity of user interfaces and screen sizes affect device and application interaction—some devices are not suitable to perform certain activities.

As a result, instead of being able to use the identity of the user as the sole deciding factor for determining access privileges, we need to take into account the user, device, application, location, and context of the

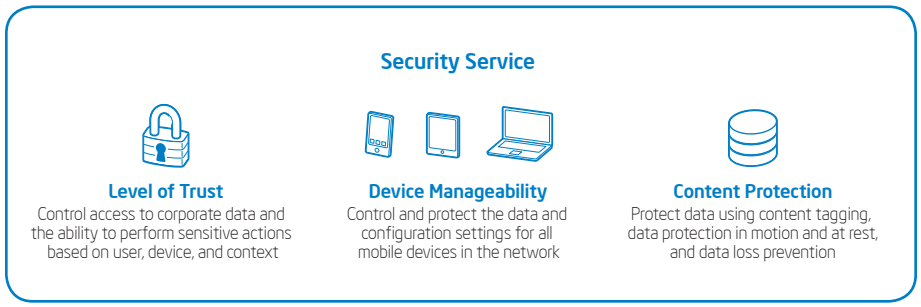

Figure 2. The security service supports a client-aware cloud by protecting the data at rest, in use, and in transit.

Intel® Architecture Supports a Client-Aware Cloud

The use of Intel architecture and solutions can improve the security of our cloud-based services.

- **Intel® Advanced Encryption Standard – New Instructions.** Accelerates encryption and decryption.
- **Intel® Identity Protection Technology with Platform Embedded Asymmetrical Token.** Provides two-factor authentication that is built into the processor.
- **One-Time Password.** This password-generating capability embedded in the platform hardware and firmware operates in isolation from the operating system.
- **Intel® Insider.** Provides an extra layer of content protection that enables secure high-definition content streaming.
- **Intel® Anti-Theft Technology.** Allows IT to remotely disable a PC if it is lost or stolen.
- **Intel® Trusted Execution Technology.** Provides a hardware-based security foundation that enables greater levels of protection for information that is stored, processed, and exchanged on the PC.
- **Intel® Cloud SSO.** Intel's identity-as-a service offering simplifies the process of providing users with access to hundreds of software-as-a-service applications, with its standards-based single sign-on, context-aware strong authentication, and account provisioning and de-provisioning.

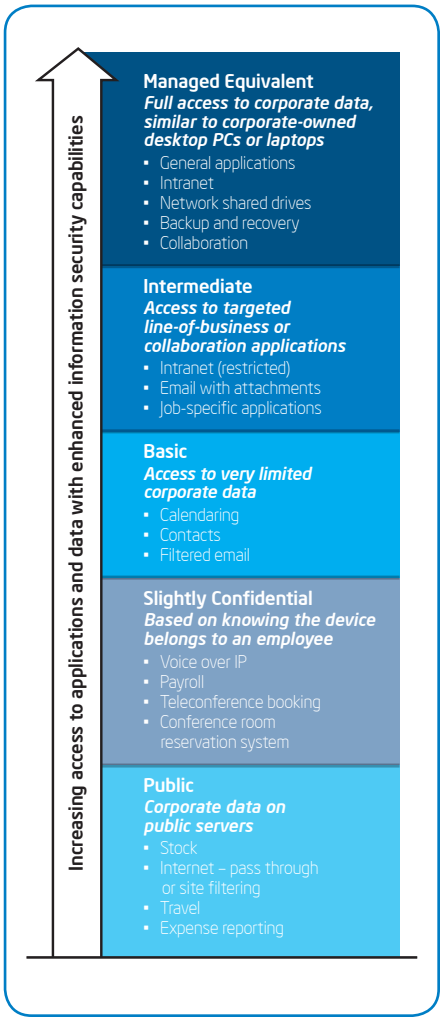

Figure 3. Varying levels of access help protect corporate data while allowing employees to use their personally owned devices at work.

task being performed. For example, we may restrict access to sensitive data in certain locations such as in a particular country, outside the factory, or off-campus.

We have developed a security model that uses technologies such as identity federation, multi-factor authentication, and certificate services for the user, device, and application, as well as location awareness, to calculate to what degree a personally-owned device can be trusted. The security model then dynamically enforces policies on that device's access by assigning it to an appropriate security level. As shown in Figure 3, each security level enables varying degrees of access and authorization to applications and data.

IMPLEMENTING DEVICE MANAGEABILITY

By controlling and protecting the data and configuration settings for all highly trusted mobile devices in the network, mobile device management (MDM) can greatly reduce support costs and business risks, enabling the secure delivery of at least a limited set of services. MDM provides a cost-effective and efficient method for system maintenance, such as replacing a corrupted or failed image with a working image. MDM also enables remote troubleshooting and the ability to remotely lock and wipe a device.

MDM is traditionally software-based. We are exploring the potential of enhancing MDM by taking advantage of hardware capabilities built into devices based on Intel® architecture (see sidebar). This will enable a higher degree of trust in the platform. For example, we will be able to assign a higher trust score to a device provisioned through our MDM solution if its credentials are protected with Intel® Identity Protection Technology.

IT cannot fully manage some devices, because the device lacks certain capabilities or features. This is especially true of personally owned and public devices. The lack of full MDM control over a device is taken into account in the calculation of trust; this affects what data and services can be accessed from the device. In some cases, this scenario leads to the use of mobile application management without the use of MDM.

ENHANCING CONTENT PROTECTION

One of the most fundamental platform components that we are developing involves a holistic approach to data protection, in a computing environment where there is an explosion of devices, OSs, and ownership models. The challenge of maintaining intellectual property protection is considerable, especially since our goal is to make data more available to improve user productivity and to protect data when the user is traveling. To accomplish this, we are exploring a combination of techniques that include content tagging, data loss prevention (DLP), and other platform- and network-aware capabilities.

A DLP solution protects data at rest, in use, and in transit. It should allow the data owner to define which data can be accessed by a user on a specific device and to control where the content can be stored. For example, in certain contexts, the DLP solution might prevent data from being copied to removable media or prevent data decryption in certain locations. Content must be protected in transit with encryption and filtered by applying policies according to file attributes, metadata, or keywords to ensure only authorized people can access it. For eDiscovery purposes, the DLP solution should also create an audit trail that tracks where content comes from and goes to.

The DLP solution is complemented by advanced protection and enforcement capabilities that include end-point security, network security and security business intelligence (BI), and a trust-aware policy framework. End-point security features include verifying system integrity, memory protection, system call monitoring, and browser security. Network security and security BI features include an advanced

sensor network, increased network access control, and a detection, remediation, and forensics environment. The trust foundation includes policy decision and enforcement, application gateways, and firewalls.

Content Synchronization Service

In order for users to have an consistent experience across multiple workspaces when using physical devices, virtual machines, and other types of containers, we need to make enterprise and personal content available to them whenever and wherever they need it, as long as such access complies with security policies. Our long-term goal is for users to be able to create and consume approved content on every approved device, although some devices may be suitable for both content creation and consumption while others are suitable only for content consumption. To attain this goal, we need to provide access to content and settings residing on multiple locations, including devices and the cloud.

If the devices have encrypted storage containers and adequate storage capacity, we may want to conditionally synchronize data between the cloud and a device, or directly between two devices. In other situations, the device may not have the capability to adequately or securely support local storage of content. In some instances, a user may specify that local storage is not desired or is desired only under certain circumstances. In these cases, we may provide access to online content without local storage.

Our content synchronization service will add business value by enabling enterprise security and scalability for content protection and sharing. This service will work with the security service DLP capabilities to protect intellectual property and enterprise data in the cloud and on the device—and to meet legal and regulatory obligations. As shown in Figure 4, to implement a fully functional content synchronization service, we need to determine what content objects to synchronize and how to organize them, decide where synchronization needs to occur, and enable the sharing of synchronized objects for collaboration.

DETERMINING WHAT TO SYNCHRONIZE AND ORGANIZING SYNCHRONIZED OBJECTS

Initially, our enterprise content synchronization service will focus on user documents, media files, and user profiles, including workspace preferences. We will also provide a service for preserving states for a consistent experience across devices. For example, the service will remember where the user left off while reviewing or updating a document.

Our goal is to synchronize content according to policy-based rules and context, as shown in Figure 4. Sample contexts include work and home, or different projects at work. To achieve synchronization, we will organize content objects into logical content groups and establish the capability to keep content groups separate. For example, we want to separate corporate and personal content groups, or separate content belonging to a specific project and make it available to others.

Over time, we will develop the ability to tag content with metadata, beyond the usual attributes that are associated with files in any file system. Examples of metadata include security classification (for example, "secret" or "public") and project lifecycle status (for example, "ratified" or "pending"). We can use the metadata to group the content to be synchronized and to make decisions about what content will be replicated to which devices and how the synchronization should occur.

DECIDING WHERE CONTENT GROUPS ARE SYNCHRONIZED

Each content group can be synchronized to one or more user workspaces. We will maintain isolation between content groups wherever they are synchronized. As stated above, we plan to accomplish this using rich metadata, but in the short term we will use a hierarchy of storage repositories and folder structures, along with more common file attributes, to define the content groups and decide where each gets synchronized.

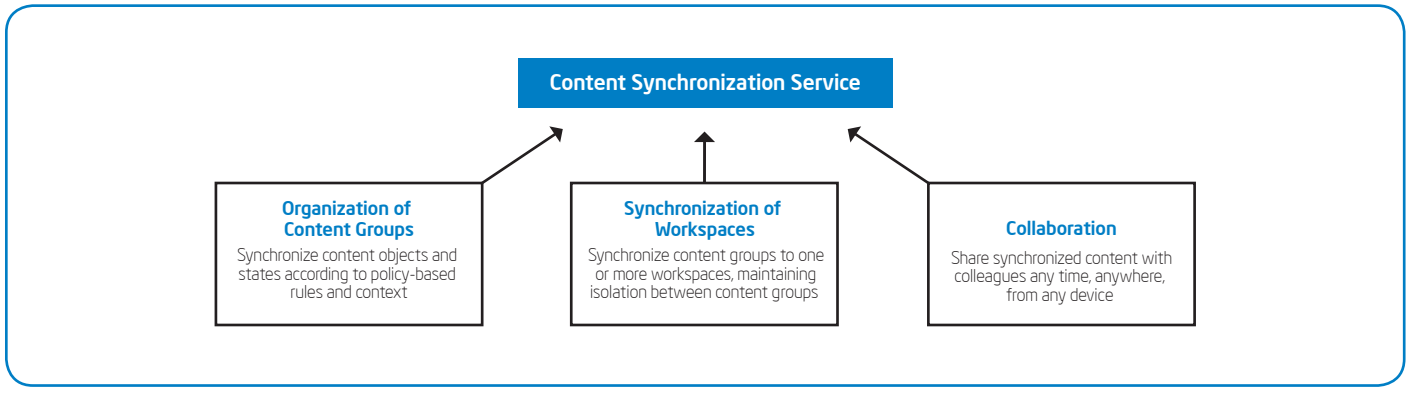

Figure 4. The content synchronization service supports a client-aware cloud by organizing synchronized objects, performing the synchronization, and enabling collaboration.

The same content object may be handled differently on different workspaces. For example, each workspace may feature different combinations of online and offline capabilities, access to public and private clouds, and encrypted or plain text storage. Some of the content objects may be automatically converted from one format to another more device-appropriate format, to suit the workspace. This is another example where a device-aware cloud service can assist in creating a more seamless user experience.

COLLABORATING ON SYNCHRONIZED CONTENT

We want to support the secure sharing of content, so that users can contact and collaborate with co-workers anytime, anywhere, and with whatever device they happen to be using. This capability will help improve employee productivity. By building a flexible infrastructure, we will be able to eventually extend the services to external contexts, such as external collaboration.

Business Application Service

Traditionally, IT has delivered services in a one-size-fits-all manner. This approach was adequate because we used one primary platform—Microsoft Windows* running on IT-built computers owned by Intel, with one browser standard. However, with the advent of many alternative compute platforms that use different OSs, form factors, and input methods, we need a more flexible approach to service delivery that can support appropriate services across a continuum of devices.

As shown in Figure 5, we envision a business application service that enables cross-platform application development, ease of context-aware connectivity between the cloud and client, and a middleware model that can provide control and translation services to devices accessing cloud-based services.

ENABLING CROSS-PLATFORM DEVELOPMENT

We are implementing an infrastructure that enables Intel application developers to write a single version of an enterprise application and allows us to deliver the service, or a subset of it, to a variety of devices. This infrastructure will optimize application delivery and provide a more consistent end-user experience across platforms.

To avoid developing applications suited for only the lowest-common-denominator devices, applications need to be client-aware. For example, an application should be able to determine how and where the device is being used and what capabilities and controls it features.

We want application developers to focus on business logic, not on underlying details. For this reason, we plan to provide a consistent platform and developer experience that as much as possible abstracts the following hardware and software differences:

- Security capabilities, such as encryption, single sign-on, and two-factor authentication, may differ from device to device.
- Access to one device's hardware components, such as the Global Positioning System, a radio, or a camera, differs from accessing similar components on other devices.
- Saving a device's state information requires different actions, depending on the platform.
- Implementing notification and push services differs between platforms.

The application platform we envision lets developers maintain consistency among different devices and states. It features a unified set of discovery capabilities that developers can use to make client-aware delivery decisions and will also re-use the security and content synchronization services.

ENABLING CONTEXT AWARENESS

To provide the best user experience, applications must be able to take advantage of platform-specific capabilities. For example, RIAs distribute processing between the cloud and the client device to improve application responsiveness. With RIAs, the code is downloaded to the client device and executes using an RIA software framework, which is typically based on HTML5. We anticipate the continuing emergence of client-aware

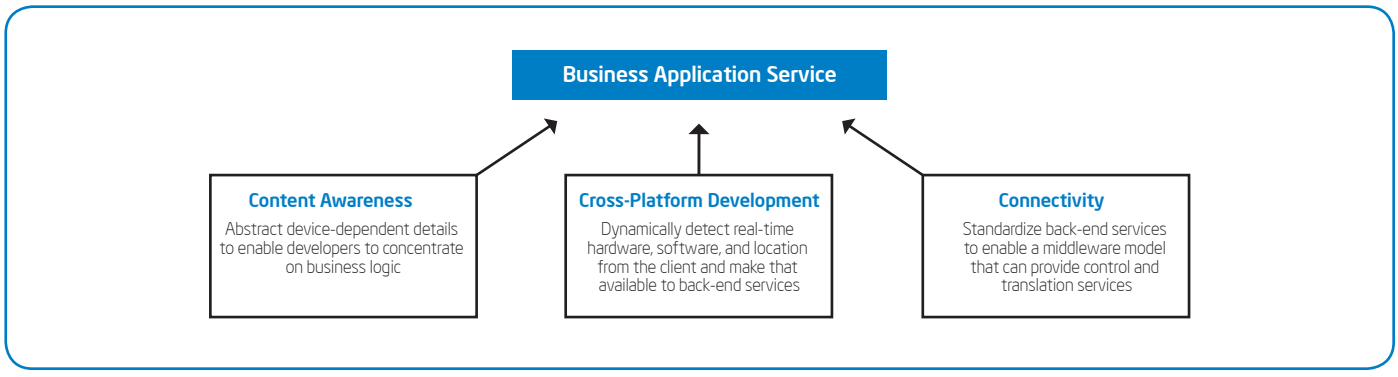

Figure 5. The business application service supports a client-aware cloud by providing capabilities for cross-platform application development, client-aware connectivity, and context awareness.

applications that can discover the capabilities of hardware, using mechanisms such as Intel® Web APIs. We are also working to enhance the performance of solutions based on HTML5 and JavaScript*.

Intel is investing in APIs that allow developers to write RIAs that detect real-time hardware information from the client, such as processor performance, battery life, and network bandwidth. Applications can use this information on a dynamic basis and remove the limitation of developing applications for the lowest common denominator. By enabling the front-end application to collect and pass context information to the back-end services, we can tailor the way we provide the back-end services to enhance the user experience .

We want to make it easy for applications to access the business services they need. To accomplish this, we are standardizing the back-end services as much as possible. For example, we will use common capabilities for licensing, metering, logging, authentication, authorization, and policies. This enables us to use a middleware model with specific servers, located in a demilitarized zone, providing control and translation services as devices attempt to access internal resources.

For example, if a smartphone attempts to access a service, the middleware server could limit the amount of information an application displays, because of the smartphone's small screen size. But if a PC accesses the same service, the middleware server will detect that a larger screen is available and display more information.

The middleware servers can also make use of situational context. For example, using physical location and Wi-Fi* network proximity provided by the client, the middleware server could determine whether a device is being used in a work or home environment and allow or deny access accordingly or prompt for additional credentials.

CONCLUSION

Our enterprise private cloud environment is now deployed across data centers worldwide, with 80 percent of new business services being deployed within the cloud. But as cloud services grow, platform diversity and consumerization of IT are increasing as well, with users expecting applications—even cloud-based ones—to be available and optimized across a wide variety of devices. We are transforming our application delivery model to enable the back-end cloud and the front-end client to work together to provide the best user experience possible. At the same time, we are aiming to improve the developer experience by making it simple to develop and deploy these robust applications across platforms.

Having previously explored client-aware cloud technologies and having developed several compelling usages associated with distributing workloads across the client and cloud, we are now implementing three foundational client-aware services that support our compute continuum and enterprise private cloud efforts.

- A security service will handle device trust, device management, and content protection.
- A content synchronization service will enable computing from anywhere by securely synchronizing content settings and states among devices and the cloud, while supporting content sharing for collaboration.
- A business application service will help developers to create secure, context-aware, and device-aware applications across platforms with access to enterprise legacy back-end services that will provide users with the robust application suite they need to be productive.

By taking advantage of the unique strengths associated with client devices and the cloud, these services provide the basis for supporting both a client-aware cloud and cloud-aware clients.

FOR MORE INFORMATION

Visit www.intel.com/it to find white papers on related topics:

- "Applying Client-Aware Technologies for Desktop Virtualization and Cloud Services"
- "Best Practices for Enabling Employee-owned Smart Phones in the Enterprise"
- "Enabling Emerging Enterprise Usages with Client-Aware Technologies"
- "The Future of Enterprise Computing: Preparing for the Compute Continuum"
- "Improving Security and Mobility for Personally Owned Devices"
- "Maintaining Information Security while Allowing Personal Hand-Held Devices in the Enterprise"
- "Rethinking Information Security"
- Information Security Protect to Enable Strategy (video)

For more information on Intel IT best practices, visit www.intel.com/it.

ACRONYMS

BI	business intelligence
DLP	data loss prevention
MDM	mobile device management
RIA	rich Internet application
SSO	single sign-on

 Printed in USA ♻ Please Recycle 0612/JGLU/KC/PDF 327471-001US