opentext™

WHITE PAPER

Data protection and PCI scope reduction for today's businesses

OpenText Secure Stateless Tokenization

Contents

"[OpenText] SST represents a paradigm shift in tokenization. It provides service at a higher performance and with greater security than con ventional, database-centric solutions. It is Coalfire's opinion that SST, when properly implemented, would effectively promote PCI compliance goals and reduce PCI audit scope for merchants and processors alike."

Coalfire
PCI DSS Scope Reduction Analysis for Voltage Secure Stateless Tokenization, October 2017

Introduction

Cyber criminals have proven adept at thwarting existing IT defenses and exploiting weak links in the payments ecosystem. Merchants, service providers, and enterprises face severe and ongoing challenges securing card data from capture through the transaction lifecycle. Organizations are also under pressure to achieve scope and cost reduction goals in meeting compliance mandates, such as the Payment Card Industry Data Security Standard (PCI DSS). Tokenization—replacing card numbers with tokens—is one method of data protection and audit scope reduction recommended by the PCI DSS Security Standards Council. Reducing audit scope using tokenization can both accelerate initial compliance efforts and reduce costs for future audits. However, merchants and service providers who have implemented first-generation tokenization approaches, whether commercial solutions or home-grown, have encountered critical problems:

- **Complexity and cost:** Managing large, replicated token databases is difficult and expensive, and these databases themselves add to PCI audit scope.
- **Integrity:** Inaccurate analytics and other application correlation due to credit card numbers sometimes being replaced by more than one token (a side effect of having a distributed token database).
- **Security and risk:** Approaches without independent and reviewable security proofs increase breach risk and do not meet QSA evidence requirements, and thus cannot achieve PCI compliance. In the event of a breach that leaks cardholder details, merchants using such approaches have no grounds to avoid significant penalties.
- **Performance and scale:** Tokenization performance is slow and very difficult to scale.
- **PCI use case completeness:** Tokenization is not suited to offline environments, such as web browsers or card swipe terminals. Supplementary solutions are required for PCI DSS audit scope reduction in such applications.

For these reasons, businesses of all sizes, from small online merchants to the largest credit card processing institutions, have found that traditional tokenization solutions ultimately failed to scale, had unexpectedly high costs, or introduced more operational issues than were acceptable in an agile and competitive environment.

This position paper is intended for merchants, payment processors, and enterprises that understand what tokenization is, and either want to replace existing first-generation implementations or implement a solution to cut costs, reduce scope, and reduce breach risks to a minimum.

Leading edge technology for tokenization and PCI DSS scope reduction

OpenText solves these critical business problems through a patent-pending breakthrough called OpenText (Voltage) Secure Stateless Tokenization (SST). This technology is part of OpenText™ Data Privacy and Protection Enterprise, an enterprise-class, payment processor- ready platform providing integrated

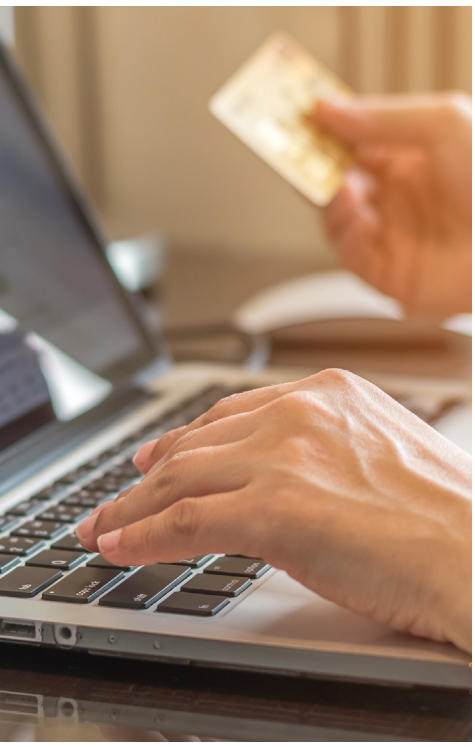

tokenization, encryption, masking, granular policy controls, stateless key management, and reporting to protect data anywhere on virtually any platform at rest, in motion, and in use. SST greatly simplifies the complexity of deploying and managing tokenization, generating consistent tokens using random numbers while completely eliminating the need for a distributed token vault or traditional database. Analytical applications such as marketing, fraud, and loyalty systems continue to operate as they did before tokenization was introduced because SST ensures that each credit card number is replaced by exactly one multi- use token. Systems based on distributed token databases cannot ensure this consistency.

SST has been designed by cryptographic experts, is based on published and proven academic research, and is validated by third-party QSAs and cryptography experts. It provides maximum protection against data exposure from security breaches while offering a proven technique for PCI DSS compliance and maximum PCI audit scope reduction.

A platform solution for compliance, data privacy, and protection

In addition to solving immediate PCI DSS compliance issues, OpenText Data Privacy and Protection is a comprehensive platform for future-proofing the compliance process. A single architectural framework combines tokenization, end-to-end encryption, structured and unstructured data protection, and data de-identification, providing data-centric security to cover a wide range of real-world data protection scenarios. Data-centric security means protecting sensitive data itself, rather than just protecting individual IT infrastructure pieces separately. Sensitive data can include payment data or any personally identifiable data as defined by PCI DSS, the General Data Protection Regulation (GDPR), Gramm-Leach-Bliley (GLB), Sarbanes-Oxley (SOX), or the Health Insurance Portability and Accountability Act (HIPAA), among others.

OpenText Data Privacy and Protection delivers tangible scope reduction for PCI DSS assessments, lower PCI compliance costs, and proven ease of deployment and management in merchant, payment processor, and enterprise environments. With leading-edge OpenText Secure Stateless Tokenization technology, it is even simpler for organizations to tokenize PCI DSS data with lower complexity and management costs than the alternatives.

Loyalty, marketing, and fraud analytic applications fail

Many types of applications depend on having a single account identifier as a primary key. If a PAN number is replaced by a token, then it must be a one-to-one replacement. Database-driven tokenization systems cannot ensure this will be the case. Standard database replication delays cause a breakdown in a variety of situations.

Example 1: Husband and wife make separate purchases at about the same time. Their cards happen to be tokenized by different data centers.

Result: Two different tokens are created for the same card number

Example 2: A single shopper makes two purchases. The master token database goes down in between.

Result: Two different tokens are created for the same card number

Limitations of traditional tokenization solutions

PCI DSS compliance is intended to reduce risk: organizations must avoid increasing risk and cost even as they aggressively pursue scope reduction goals. Qualified security assessors (QSAs) recommend tokenization of PAN data to reduce PCI scope, and the PCI Security Standards Council has issued guidance on how to achieve it, but not all tokenization systems are created equal. Most tokenization systems suffer from some or all of the following:

- An expensive and cumbersome database management problem
- A requirement for a "token vault," which actually adds to PCI audit scope
- An unproven, probably insecure token generation method

Traditional tokenization involves building a large-scale database ("token vault"), storing each PAN together with a randomly generated token. The size of such a database makes it impossible to pre-tokenize all possible credit card numbers, so PANs are added as they are received in transactions, and the table grows over time. Each PAN/token pair must be replicated to multiple copies of the database to meet business service reliability. This leads to the following critical problems:

- High cost of licenses for database and replication software
- Skills requirement for database administrators to handle complex replication and recovery scenarios
- Advance purchase of sufficient database and storage hardware to handle seasonal peaks, such as "Black Friday"
- Key rollover in the token vault—typically a lengthy, resource-intensive, manual process
- Variable one-to-one correspondence between card number and token: replication delays can mean that sometimes two tokens are generated for the same card number, resulting in failures or inaccuracies in applications such as loyalty, marketing, and fraud analytics

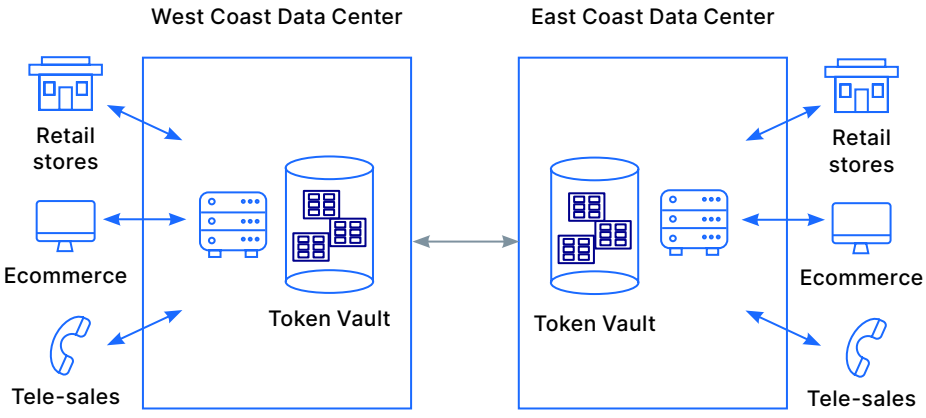

Figure 1. Ideal data path for traditional protection approach

OpenText Secure Stateless Tokenization

OpenText Data Privacy and Protection Enterprise includes a new technology for generating tokens without the high costs and application failures inherent in traditional approaches that use an ever-growing database. Secure Stateless Tokenization uses a set of static, pre-generated tables containing random numbers created using a FIPS random number generator and based on published and proven academic research. These static tables reside on the OpenText Data Privacy and Protection Enterprise appliances and are used to consistently produce a unique, random token for each clear text PAN input, resulting in a token that has no relationship to the original PAN. No "token vault" database is required, thus improving the speed, scalability, security, and manageability of the tokenization process.

This offers a number of advantages:

- No external database software or hardware acquisition or management costs
- Greatly simplified management with lower maintenance costs:
    - No database growth over time—No old PAN/token mappings to be purged over time
    - No replication and backup issues
- 100% consistency in the PAN/token mappings provided by all servers in all data centers
- Always a one-to-one correspondence between PAN and token: two tokens never exist for one PAN
- High performance: tables are small enough to fit in memory, enabling high-speed tokenization
- Key rollover (changing the encryption key that protects the random number tables) takes minutes
- Improved security: the randomly generated tables do not allow an attacker to determine which tokens are in use or which PANs have recently been tokenized since SST does not require storing a list of credit card numbers as part of its lookup tables.

The static tables consist of random numbers, are unique to each OpenText Data Privacy and Protection Enterprise installation, never need to be changed, and do not grow over time. The static tables are copied securely to all OpenText Data Privacy and Protection Enterprise servers where tokenization will occur. As a result, any number of OpenText Data Privacy and Protection Enterprise appliances can be deployed for redundancy or for high throughput. The result is that enterprises can:

- Maximize PCI DSS audit scope reduction.
- Drastically reduce tokenization deployment and management costs, while at the same time improving security and scalability.

Global Credit Card Processor

Business drivers:

- Offer value-added PCI scope reduction services in e-commerce and tokenization

Situation:

- Stratus VOS, UX, and NonStop across multiple data centers
- Solution must have one-to-one correspondence between token and card numbers
- In-house database-driven tokenization couldn't scale, limited performance, too complex to ensure token consistency

Solution:

- OpenText Data Privacy and Protection Enterprise with OpenText SST and OpenText Data Privacy and Protection Enterprise Web PoC completed quickly

- Keep analytical applications running smoothly for business applications such as loyalty, marketing, and fraud.

This technology can also be provided as a cloud service, so organizations who do not want an on-premises solution have the option of simple tokenization as-a-service with the power of total ownership and control.

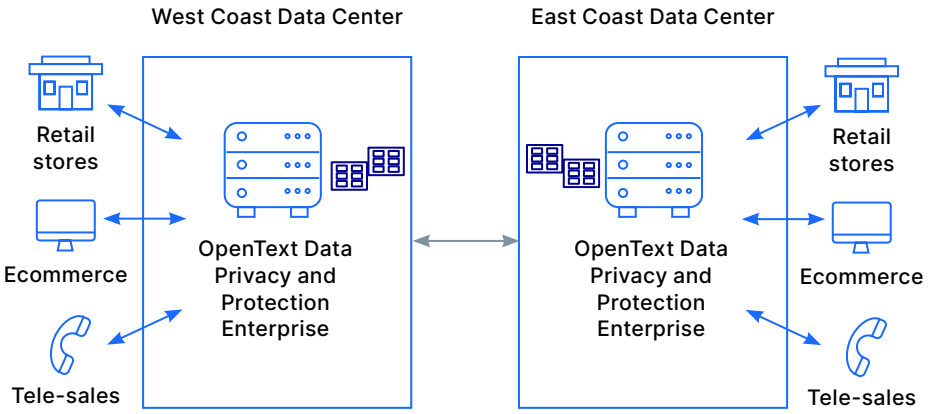

Figure 2. OpenText Data Privacy and Protection Enterprise tokenization: one-time distribution of static token tables

Token multiplexing

Security best practices and PCI DSS guidance both point to the need to make tokens meaningful and usable only to the particular group of applications that require them. This PCI DSS guidance states that if this is not the case, then these "high-value tokens" may still be in scope.

This adds tokenization requirements beyond PCI scope reduction:

- Service providers supporting multiple merchants need a simple, scalable way to provide a different set of tokens to each of their clients
- Individual merchants and other organizations need a means to provide a distinct set of tokens to each independent business unit or application group

To meet these requirements, first-generation tokenization products force organizations to implement separate databases—or even separate tokenization systems—for each line of business or merchant, to ensure tokens are not exposed across environments. This is not only extremely costly but presents tremendous scalability challenges.

OpenText Data Privacy and Protection Enterprise is unique in supporting token multiplexing: a simple way to create token independence between merchants, applications, or line of business, avoiding the cost and complexity of multiple databases. Powered by the combination of randomly generated tokens and OpenText Format-Preserving Encryption, this is a "zero provisioning" feature—there is no requirement to configure OpenText Data Privacy and Protection Enterprise separately for each new merchant or business unit using tokenization. Randomly generated tokens are converted to merchant- or

US Airline

Business drivers:

- Failed PCI audit, facing steep fines
- Remove e-commerce web servers from PCI scope
- Protect card numbers from mobile devices during the flight and ACH payment data

Situation:

- Support for e-commerce, call center, mobile, inflight, and retail
- Multiple programming languages, data in Microsoft SQL Server

Solution:

- OpenText Data Privacy and Protection Enterprise with OpenText SST brought systems into PCI compliance, took e-commerce web servers and mobile payment data out of PCI scope

business unit-unique tokens using a specific Format-Preserving Encryption (FPE) policy, isolating tokens to the specific entity receiving them. FPE is an additional step performed on the randomly generated tokens, which makes the tokens unique without altering their format—for example, a 16-digit card number with a valid Luhn check digit will remain 16 digits with a valid Luhn check digit, if that is the desired format of the token.

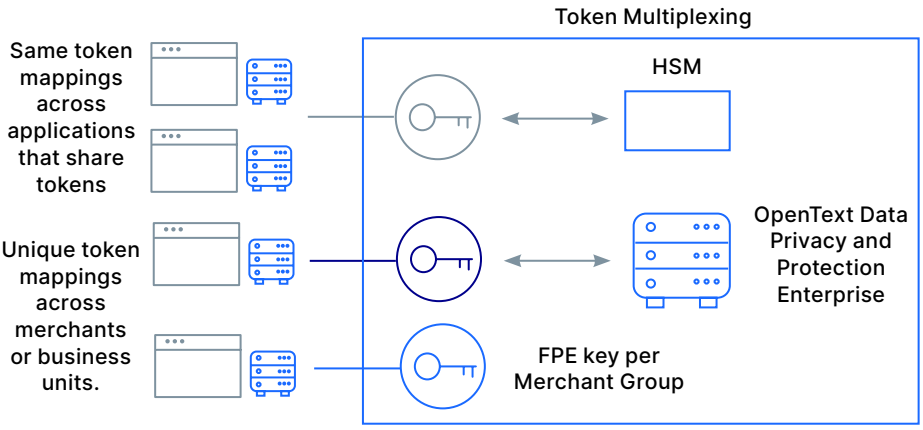

Figure 3. OpenText Data Privacy and Protection Enterprise Token Multiplexing simplifies token independence between merchants, applications, and line of business, avoiding the hassle of using multiple databases

Beyond tokenization for PCI DSS scope reduction

There are some areas of infrastructure where PCI data needs protection and PCI DSS audit scope can be reduced, but tokenization is not practical. For example, environments such as point-of-sale (POS) terminals may be offline for substantial periods of time and therefore cannot depend on a centrally managed tokenization service. Security considerations mean that other environments, such as consumers' web browsers, must not have direct access to tokenize, either through external APIs or by having a local copy of a token lookup table.

OpenText Data Privacy and Protection Enterprise offers techniques to protect card data immediately after capture, minimizing PCI audit scope in those environments as well as in the areas where tokenization applies. These points of capture include:

- **POS terminals**— OpenText Data Privacy and Protection Payments encrypts data on the terminal hardware itself; data is translated to random tokens once received by the back office
- **Consumer web browsers**— OpenText Data Privacy and Protection Enterprise Web encrypts data on the web browser before it gets sent to a merchant's web servers. When properly implemented, this drastically reduces the PCI scope of the web servers, while preserving the merchant's control over every step of the payments process

All OpenText Data Privacy and Protection Enterprise solutions share a common infrastructure, including the same centralized servers and administration tools. This enables OpenText Data Privacy and Protection Enterprise customers

to choose an appropriate combination of techniques to maximize PCI scope reduction across diverse environments while avoiding the costs of deploying and managing multiple products.

For more information on how OpenText Data Privacy and Protection Enterprise solutions can meet PCI DSS requirements, see the OpenText Privacy and Protection Foundation product overview.

OpenText Data Privacy and Protection architecture

Virtual appliance

OpenText Data Privacy and Protection Enterprise is delivered as a virtual appliance. The package installs directly in a standard Intel-based physical server or on a virtual machine and can be up and serving tokens within a few hours. This appliance is all-inclusive, with no requirement to install and maintain separate software components such as application servers and databases. Applications, network devices, and databases can tokenize and detokenize using the OpenText Data Privacy and Protection Enterprise web services API.

Optional client components

Optional client-side software components enable tokenization to be performed using local API calls or command-line operations and can be scripted for high-throughput batch operations .

Rich formatting options

PCI DSS allows certain digits of a credit card number to remain “in the clear”—normally the first six digits or the last four. OpenText Data Privacy and Protection Enterprise tokens optionally preserve these digits, which means many business processes can work with tokens without the need to de-tokenize.

Tokens can also be configured to exactly match the format of the original data, such as a credit card number (between 12 and 19 digits), or to substitute alpha characters for certain digits. This enables tokens to be immediately identified as tokens, which can simplify audits .

These formatting settings are centrally controlled through an easy-to-use web interface and enforced as policies across all tokenization servers and APIs.

Identity management, authentication, and authorization

OpenText Data Privacy and Protection Enterprise supports tokenization best practices in system management, change control, and operational administration by enforcing authentication and authorization on critical operations that may compromise the security of the data protection system. The system can use LDAP, Active Directory, or other credential stores to authenticate both users and applications.

Identity authorization rules require a specific authentication match using a choice of methods set up by the administrator. These methods can be used individually or layered for added security.

Medium-sized e-commerce electronics retailer

Business drivers	Situation	Solution
Not compliant due to new requirements in PCI DSS, need to get compliant quickly!	Current security controls left systems in PCI scope and noncompliant	OpenText Data Privacy and Protection Enterprise with OpenText SST moved all systems from PCI scope, with higher security and audit scope reduction compared to other solutions
Minimize audit scope and PCI footprint for simple compliance process	All Microsoft environments with three major applications handling tokens: POS, contact center, and e-commerce	Up and running in one day, with annual key rotation taking only five minutes Easily integrated with their all Microsoft (SQL Server, .NET) environment

Fine-grained permissions

To provide the ability to lock down access to sensitive data as much as possible without impeding business processes, OpenText Data Privacy and Protection Enterprise can limit access by individual applications or users in the following ways:

- Tokenize or de-tokenize only
- De-tokenize but with certain digits “masked” (replaced by asterisks or other dummy characters)

Dual controls

Just as critical to tokenization best practices is enforcement of dual controls to ensure a malicious administrator cannot independently compromise the security of the tokenization system. Unlike systems that introduce cumbersome processes for dual controls, OpenText Data Privacy and Protection Enterprise supports workflow-based dual controls so administrative tasks are still easy to set up and execute. The following are some examples of administrative functions that can be configured for dual controls in OpenText Data Privacy and Protection Enterprise:

- Adding and editing authentication methods
- Editing system-level data protection settings
- Adding and editing administrators
- Backing up and restoring the console configuration

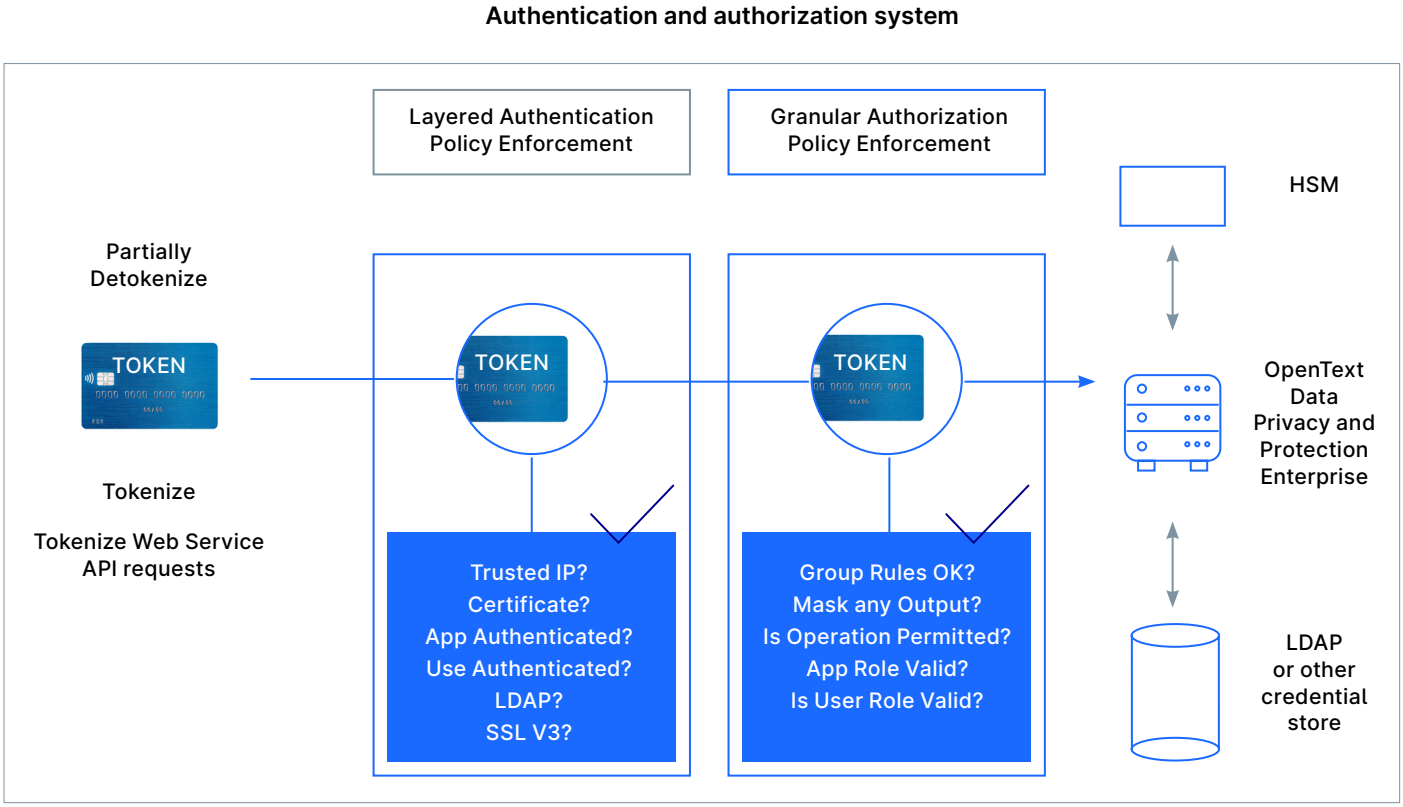

Figure 4. OpenText Data Privacy and Protection Enterprise supports tokenization best practices, authenticating both users and applications

Conclusion

The OpenText Data Privacy and Protection Enterprise platform with OpenText Secure Stateless Tokenization offers a superior solution for merchants, payment processors, and enterprises to achieve compliance, reduce PCI DSS scope to the maximum extent possible, and protect cardholder data—all without the complexity, cost, and scalability issues of first generation tokenization solutions. The OpenText Data Privacy and Protection Enterprise platform provides industry best practices support for a diverse range of compliance requirements, including legacy and contemporary systems, at the scale and performance level required by merchants, payment processors, and enterprises. Its unified administration, simplified key management, ease of ongoing maintenance and compliance ensures that an investment in OpenText Data Privacy and Protection Enterprise today will continue to meet future security and compliance needs—from mainframe to mobile, end to end.

Learn more