GridGain

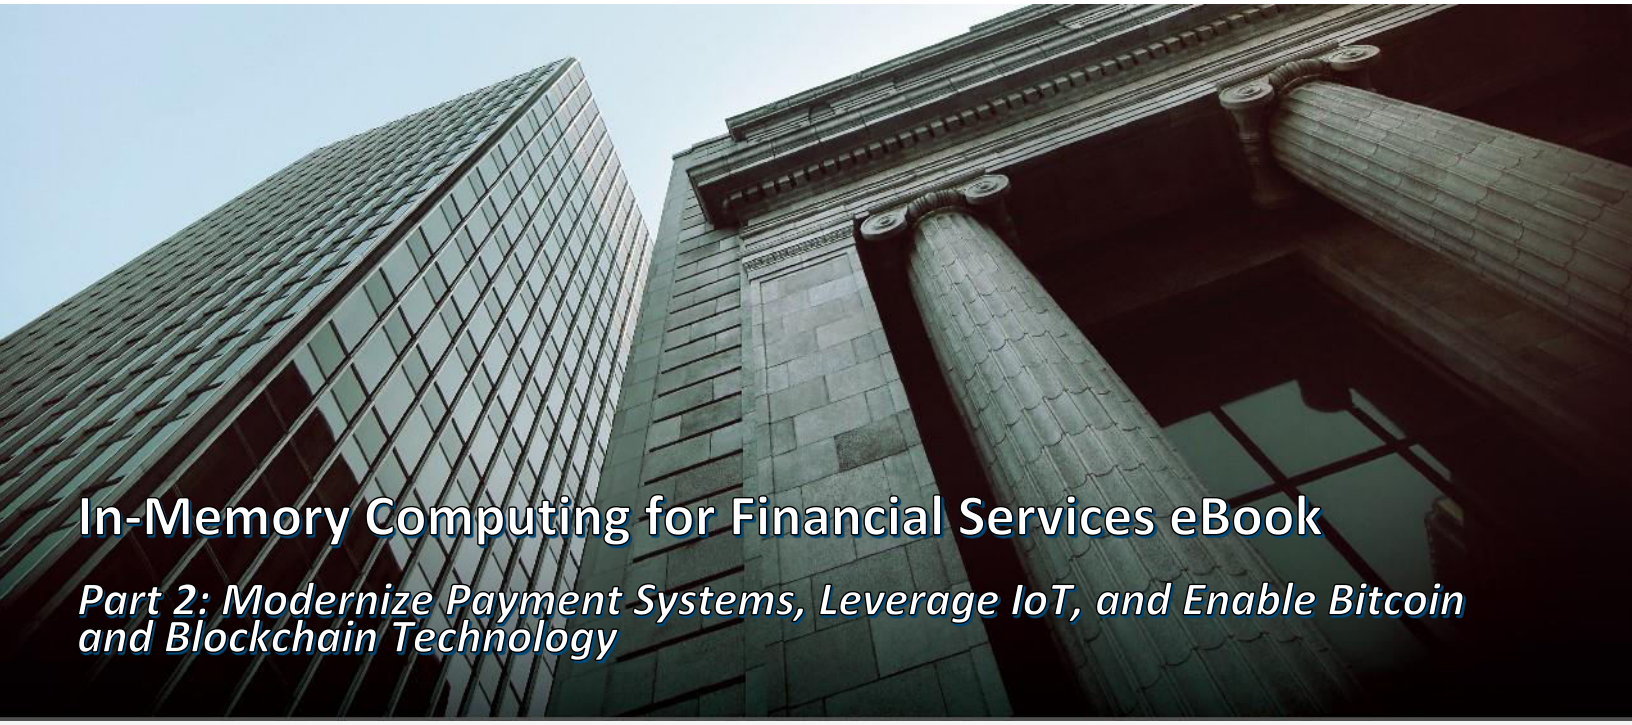

In-Memory Computing for Financial Services eBook

Part 2: Modernize Payment Systems, Leverage IoT, and Enable Bitcoin and Blockchain Technology

A GridGain Systems In-Memory Computing eBook

October 2017

Table of Contents

Chapter 1

Modernize and Accelerate Payment Solutions with In-Memory Computing

The payments industry is evolving rapidly. The shift to digital payments is taking place on many fronts, from bitcoins and mobile wallets to "tap and go" payment transactions and peer-to-peer money-transfer apps. According to Business.com, fewer than 25 percent of U.S. in-store purchases in 2017 are expected to be cash-based transactions. Worldwide, the mobile payments market has grown from $235 billion total revenue in 2013 to projected values of almost $800 billion in 2017 and over a trillion dollars by 2019. Digital and mobile technologies are radically disrupting the ways that payments are initiated and processed – and increasing the potential to mine payments data for actionable insights.

As providers of payment services seek to stand out in an increasingly crowded field, technology is a crucial part of maintaining a competitive edge. With consumers expecting instant, invisible payments and money transfers, providers must be able to process payment transactions at real-time speeds.

They must also be able to scale their capacity quickly to meet usage growth. Plus, they need sophisticated analytics capabilities to help them prevent fraud, handle regulatory compliance, and identify revenue- generating insights from the plethora of data they are collecting.

To achieve this level of performance, scalability, and analytical sophistication, many payments providers are turning to in-memory computing solutions.

This chapter will discuss changes in the payments industry, the new opportunities and challenges providers are facing, and how providers can gain the edge they need with solutions such as the GridGain in-memory computing platform.

How the Payments Industry is Changing

The payments industry has undergone radical changes in recent years, both in visible ways and behind the scenes.

Key changes include the following:

- New digital & mobile payment methods
- Efficient new methods for tracking payments
- Heightened security measures
- Use of data-driven insights to increase profits

Let’s take a closer look at these changes and how they are presenting new opportunities and challenges for payments providers.

New Ways to Pay

Almost every type of payment situation has been transformed by new methods of making payments digitally. Whether it's consumers paying for goods and services or people sending money to each other, the way money flows from one party to another is looking much different now than it did a few years ago – creating new opportunities and challenges.

Trends and opportunities. New trends in payment methods include the following:

- Instant, invisible payments from digital devices
- Mobile phones and wearables increasingly used as payment devices
- Mobile wallet apps, such as ApplePay, Android Pay, Samsung Pay, Walmart Pay, and Bitcoin gaining wider acceptance
- Embedded payment options allowing payment from within websites and apps, including social media
- Contactless "tap and go" spending via chip-enabled cards (such as transit passes) or NFC-enabled smartphones that communicate with RFID or NFC readers
- Faster ACH payments (through the Automated Clearing House Network in the U.S.) – same-day, in many cases – of payroll direct deposits, mortgage payments, and other authorized transfers
- Peer-to-peer digital payments (remittance) – including across borders and currencies – via apps, instead of through Money Transmitter Operators (MTOs)
- IoT devices initiating payments – for example, an air conditioner automatically requesting air-filter service and paying with credit card information

Innovative trends such as those noted above – and the lower cost of digital technologies in general, compared with cash payments – are creating significant new opportunities for businesses. Many technology firms are specializing in payments and attracting funding from venture capitalists who strongly believe in the future of their technologies. Traditional financial firms are also partnering with fintech firms in ventures such as mobility innovation labs.

In addition to being technology driven, much of the opportunity in the payments industry centers around creating a better customer experience. Banks and merchants can gain business by making it easier and faster for customers to access their money (or their credit) and enabling them to make payments from all of their devices and apps. New entrants to these markets can also gain a lot of business by being cheaper or providing better service than more established competitors – particularly in areas such as digital remittance, formerly dominated by MTOs such as Western Union and Money Gram.

Challenges for providers. Keeping up with the new trends and staying ahead of the curve on customer experience both pose challenges for payments providers. These challenges include the following:

- Offering payment options from a range of devices and apps – including mobile and wearable devices, social media, and other options mentioned under "Trends and opportunities," above – while maintaining a consistent customer experience
- Easing culture change to help older customers get on board with new ways of making payments

- Maintaining high performance and 24/7 availability wherever possible, in order to meet user expectations of instantaneous payments
- Accessing real-time data on currency conversion rates – and rates that competitors are offering for payments made across currencies
- Accommodating customer growth through a scalable infrastructure that keeps performance and service from degrading as the number of customers increases
- Cooperating with other players, as new scenarios demand more collaboration among vendors, banks, and Fintech

To meet these challenges successfully, providers need technology that meets a critical set of criteria: it should be interoperable, high performance, low latency, and scalable – ideally open source and cloud-based. For payment situations requiring real-time analysis of large datasets, such as current currency conversion rates for international payments, the technology should also be able to handle big-data analytics. We'll discuss these technology needs in more depth after we look at some of the other trends and challenges in the payments realm.

Turbo-Charged Payment Tracking

Behind the obvious changes affecting how people make payments, significant changes are also occurring in how payments are processed and recorded. Multiple trends are creating the potential for a much faster, more efficient process.

Trends and opportunities. Some key trends in tracking of payments include the following:

- Decentralized, distributed public ledgers based on the blockchain technology behind Bitcoin and able to greatly streamline transaction processing
- Payment hubs, which replace separate, specialized payment services with a unifying central infrastructure that provides a common set of services integrating all payment applications
- Reduced overhead and improved efficiency as a result of both of the above trends

Providers who take advantage of the efficiency improvements available through payment hubs and decentralized public ledgers, or blockchains, can gain a strong competitive advantage through faster, more streamlined processing and lower costs. *(For an in-depth discussion of how blockchains can potentially streamline processing even for non-bitcoin transactions, see the Chapter 3, "Enable Bitcoin and Blockchain Technology with In-Memory Computing.")* However, these opportunities are accompanied by some challenges.

Challenges for providers. As the pace of digital payments accelerates, providers face several challenges:

- Processing high volumes of streaming data – particularly with blockchain-based systems, since there is no central data repository; blockchain subscribers must quickly process and store incoming data
- Maintaining high performance as data volumes grow
- Providing high reliability and availability at the level required for financial systems

High-performance, scalable, fault-tolerant data systems with streaming capabilities, high availability, and recoverable transactions are needed to meet these challenges. These are some of the same criteria needed to address another important trend in the payments industry: the increase in security risks and regulatory requirements.

Heightened Security Measures

As online, mobile, and cashless transactions fuel the growth in digital payments, the potential for fraud, hacking, and other security risks escalates. Those who want to commit fraud now have ways to do so more easily and on a much larger scale than ever before, as text dumps of credit card information turn into gold mines for enterprising cybercriminals.

Much attention is now focused on figuring out how to protect transaction data and how to detect and prevent fraud quickly enough to avert harm. Increased cybersecurity regulations are promoting greater security – but also increasing the challenges for payments providers who must implement them.

Trends and opportunities. Trends in payment security include the following:

- Increased adoption of card-securing EMV solutions such as chip-enabled credit and debit cards
- Increased payment safety requirements in PSD2, the EU's revised Payment Services Directive, including requirements for stronger identity checks in online payments
- More countries promoting cybersecurity initiatives, both individually and as part of regional bodies, such as initiatives from the Organization of American States and Europe's Organization for Economic Co-Operation and Development
- Focus on encryption requirements for NFC (near field communications) to prevent information theft during contactless payments
- Focus on secure authentication and authorization through techniques such as biometrics and increased requirements for digital signatures

Providers who can successfully navigate these trends and maintain security and privacy for their customers can gain a significant edge in reputational integrity. As PayPal CEO Dan Schulman noted in a recent radio interview, "When it comes to financial services, the single most important brand attribute you can have is trust."

Preventing fraud and hacking also decreases financial loss and improves a company's financial competitiveness. However, meeting fraud-prevention goals can be quite challenging.

Challenges for providers. Maintaining secure and compliant payment transactions involves significant challenges, including the following:

- Implementing effective security and encryption techniques for user authentication and protection of transaction data
- Understanding and complying with security regulations – often highly complex ones – and doing so in real time

- Accessing large amounts of relevant data – both structured and unstructured – such as historical behavior data that can help predict future fraud
- Identifying and stopping fraudulent transactions in real time through predictive modeling, machine learning, and other techniques for flagging suspicious behavior

Making effective use of data is crucial in helping providers to meet these challenges. They need systems that can analyze massive amounts of data, so they can leverage everything they can find out about customer identity and historical behavior and then use sophisticated algorithms to help them spot behavior that is abnormal and potentially fraudulent. They also need these types of big-data analytics to help them comply with increasingly rigorous security regulations. Plus, they need all this to happen at real-time speeds, to stop fraud in its tracks.

(For more detail on processing requirements for fraud detection and regulatory compliance, see the GridGain white papers "Powering Financial Fraud Detection and Prevention with In-Memory Computing" and "Achieving Real-Time Financial Regulatory Compliance with In-Memory Computing.")

The same fast-data analytics capabilities that help providers maintain security and prevent fraud can also be a boon with respect to another emerging trend in the payments industry: mining payments data for "actionable insights" that can increase profits.

Profitable Data-Driven Insights

As data mining becomes an increasingly important source of profit for companies, payments providers are uniquely situated to benefit. They have access to payment data not just for one company, but across a wide range of vendors, services, and payment recipients. They may also have access to other illuminating information about customers – such as the comments with which many users tag their payments when they use Venmo to transfer money to others.

With this unprecedented degree of access, payment-services providers are well positioned to use their data to generate insights that can drive new business opportunities.

Trends and opportunities. Trends in leveraging payments data include the following:

- Combining payments data with other data sources such as geolocation data, customer history data, and social media information that can provide insights into what is going on in a customer's mind
- Incorporating better and faster analytics to hone in on a customer's state of mind, predict possible future customer behavior, and identify "actionable insights" – that is, opportunities to increase profits or customer satisfaction – in real time

Using payments data to identify actionable insights in real time can provide tremendous strategic advantages. Providers can use payments data to identify in-the-moment opportunities for selling a customer additional products or services (upselling), selling related products or services (cross-selling), or taking actions that can lead to improved customer retention.

Challenges for providers. As with preventing fraud, generating actionable insights from payments data is a knowledge-intense activity that must be performed at real-time speeds to be effective. It requires accessing large amounts of data – including potentially unstructured types of data from a variety of sources – and analyzing it quickly enough to identify and seize business opportunities in the moment of time when they are relevant. For these reasons, a high-performance system capable of fast-data analytics, machine learning, and other A.I. techniques is essential.

Meeting Payment Challenges with State-of-the-Art Data Technology

In the new world of digital and mobile payments that must be tracked with high efficiency and tight security – and that can be profitably mined for data-driven insights – providers need a state-of-the-art data technology to meet their new challenges successfully.

This data solution must be above all fast and scalable, as well as secure, reliable, flexible, and capable of sophisticated big-data analytics. In other words, it should have the following characteristics:

- High-performance and low latency: Speed is essential for providing the instant payment services that customers expect, processing high volumes of streaming transaction data, and providing the real-time analysis needed to identify both potential fraud and actionable insights for immediate business opportunities.
- Scalable: The ability to easily scale a system up is essential for accommodating fast growth and maintaining high performance even as payment volumes increase.
- Flexible and interoperable: Flexibility (for example, being open source and cloud-capable) and interoperability are important for offering payment options from a range of devices, coordinating payment services via hubs, and enabling access to diverse data sources for fraud prevention and greater customer insight.
- Fault-tolerant and highly available: Payments data must be maintained with the high level of reliability, availability, and recoverability required for all financial data.
- Secure: Fully implemented security features are vital for keeping customer financial data safe.
- Capable of sophisticated big-data analytics: The ability to effectively use predictive modeling, complex event processing (CEP), machine learning, and statistical and A.I. techniques is extremely important for supporting fraud prevention and data-driven business intelligence in real time.

Fortunately, there is a data technology that is available with these features and is well suited to this type of high-speed, big-data use case: in-memory computing, as implemented in the GridGain in-memory data platform.

Chapter 2

Enable Bitcoin and Blockchain Technology with In-Memory Computing

Bitcoin and blockchain, the digital-ledger technology behind this electronic currency, are generating enormous amounts of interest in the financial services industry. Most of the larger banks are investigating this area, and many technology companies are building platforms to enable blockchain technology for financial services firms.

In a 2016 Deutsche Bank survey of 200 participants in the global financial industry, almost all respondents (87 percent) expected this revolutionary technology to have a major impact on the securities services market. Three-quarters of the respondents said they expected to see widespread adoption of blockchain technology within the next three to six years. In addition, almost two-thirds expected blockchain adoption to produce substantial savings, in the range of 11 to 25 percent.

However, the adoption of blockchain technology will require significant technology upgrades. Current legacy systems cannot provide the security and performance needed to support the decentralized blockchain architecture in a well-integrated way. Participants will need scalable, distributed, and well-protected computing platforms in order to leverage the technology successfully.

This chapter will discuss bitcoin and blockchain technology, describe the innovative opportunities this technology offers for financial services firms, and examine how in-memory options such as the GridGain in-memory computing platform can address the challenges and boost the effectiveness of tomorrow's blockchain applications.

Before diving into the details of bitcoin and blockchain, let's look at some of the issues behind the need for these technologies.

Issues in Today's Financial Environment

The widespread interest in blockchain technology springs from its potential to help address several current issues with the financial ecosystem, including:

• **Lack of trust**. With recent volatility around banks coming into and disappearing from the financial ecosystem, there is a great deal of uncertainty and mistrust. People and institutions are not certain who can be trusted with their money and their payments. In addition, the growing problem of sophisticated financial fraud has created doubt about whether transactions come from a reliable source, increasing the mistrust.

• **Isolated legacy systems and processes**. Many financial services companies today are using legacy systems that have trouble connecting with other systems, creating isolated data islands.

• **Too many intermediaries**. Today's transaction processes are a complex series of steps. Some of these steps involve overnight or batch-oriented processing. This increases the time needed to get through the process as a whole.

• **Slow and complex post-trade functions**. Automation initiatives in the trading business, such as straight-through processing, have sped up the trading process itself. However, in the post-trade space, high complexity and numerous intermediaries continue to slow down processing.

• **Time-consuming reconciliation**. Within a financial system, everything must work smoothly, with no loss of records. In a sequential series of processes, there must be reconciliation after each process to ensure that all transactions are accounted for. And, of course, reconciliation involves numerous record-keeping, record-checking, and record-enrichment steps – all of which create additional complexity.

• **Time-consuming compliance checks**. Since the financial crisis of 2008 through 2010, the financial services environment has been heavily regulated. All transactions must be checked for regulatory compliance, with rules differing depending on whether the environment is banking, trading, or payments. These compliance checks add numerous validation and rules-checking steps to each process.

Before we explore how blockchain technology can help to overcome these issues of mistrust, data isolation, and complexity of processing, let's look at what blockchain technology is and how it is currently being used – the most prominent example being the role of this technology in bitcoin.

Bitcoin: A New Model for Electronic Currency

Bitcoin is a cryptocurrency – a fully electronic currency secured by cryptography – that operates without oversight from any central bank authority. Originally envisioned as a framework to allow small payments to be made electronically without incurring credit card fees, bitcoin technology has become much more. Bitcoin with a capital "B" is now a flourishing protocol and a payment network that is used in many areas within the general category of payments. There are virtual exchanges for buying and selling bitcoins (which are stored in electronic "wallets"), as well as for exchanging them with traditional currency, and bitcoins have performed well as investments. According to Bloomberg, their value has surged more than 4,000% since 2011.

Where bitcoins come from. Bitcoins are introduced into circulation using a mechanism called *mining*. Mining involves competing with other parties to perform transaction-recording tasks and obtain bitcoins as a reward. This reward creates the incentive for miners throughout the internet (anyone with sufficient knowledge and processing power) to perform the record-keeping tasks that are necessary to perpetuate the system. Miners record transactions in a public, distributed ledger called the blockchain. (Note: The word "Blockchain," when capitalized, can also refer to a software and wallet company that hosts a real-time browser of bitcoin transactions. Throughout this paper, we use the word to refer to digital ledgers and ledger technology in general.)

Challenges to bitcoin adoption. Some potential adopters are approaching bitcoins cautiously due to a variety of issues. Because bitcoin is a fully electronic currency, it requires special setup steps not needed for familiar currencies, and it is more vulnerable to cyber-theft, hacking, authentication issues, and

system failures. Because it operates outside of existing currency controls, it is more volatile than other currencies and is also the focus of increasing government scrutiny. In addition, there are concerns about system independence, since about 70% of the bitcoin mining is currently done in China, according to Bloomberg. As familiarity with bitcoin technology grows, many or all of these concerns are likely to ease.

Wider implications of the bitcoin model. Regardless of the degree to which bitcoins are embraced as a currency, there is growing interest in the digital-ledger technology behind bitcoins – that is, the blockchain. This technology offers substantial promise for streamlining many types of transactions.

Blockchain: The Revolutionary Technology Behind Bitcoins

Describing a blockchain simply as a public, distributed ledger of transactions does not fully convey its power and potential. What makes blockchain technology particularly compelling is that it creates a decentralized, self-perpetuating infrastructure that transmits information simultaneously to all interested parties. This model can save substantial amounts of time and money for those needing access to the transaction data.

How blockchains operate. Whenever a transaction occurs, such as a person sending bitcoins to a company to purchase a product, a secure transaction message goes from sender to receiver (say, person to company). Only the sender and receiver can view the contents of the message, but a record of the transaction (that is, of the particular bitcoins being spent) is added to the blockchain and is viewable while encrypted by all blockchain subscribers, who can see the source and destination. If they need to see the contents, they decrypt the block.

Direct access to transaction data. Because all subscribers receive the blockchain update simultaneously, there's no need for costly intermediaries or a time-consuming set of reconciliation steps to complete the transaction.

Transparency plus security. While it might seem as though such a public infrastructure would be vulnerable to tampering and errors, the design of the blockchain ensures a very high level of integrity and security. A blockchain is a set of sequential blocks of encrypted transaction information that can never be removed or deleted, only added to, in chronological order. There is one blockchain for all bitcoin transactions, and each block of transactions contains both a timestamp and a hash of the previous block, making the blockchain highly tamper-resistant.

The randomness inherent in transaction verification also contributes to the security of blockchains. Each transaction is sent randomly to one of many computers that subscribe to the blockchain, and that computer can then verify its validity (in terms of source and destination, plus contents if it is a bitcoin transaction) before distributing it. The awareness that anyone anywhere might end up validating the transaction helps to deter those who might want to interfere with transactions; they can't influence the validator because there is no way to predict who the validator will be.

An open, decentralized model. The blockchain model is based on decentralized, peer-to-peer transactions. All subscribing computers can communicate with the blockchain and view records within an open, cloud-like interface – sometimes referred to as the "Internet of Trust."

This idea of direct connection between involved parties without much centralization is key to understanding blockchain technology and why it matters to the financial community. Transaction-processing systems based on this direct, decentralized system have tremendous potential to overcome the mistrust, isolation, and complexity that have plagued the financial ecosystem.

Potential Blockchain Use Cases

Many banks, large financial institutions, registry houses, custodians, and payment centers are evaluating blockchain technology. Because all blockchain subscribers receive direct, simultaneous access to transaction data, few counterparties and transaction checks are needed. Each transaction requires only a couple of records to be sent, offering substantial streamlining compared to current transaction processes.

Some potential blockchain use cases for financial services include the following:

• **Post-trade processing**. Back-office post-trade functions such as netting, matching, and clearing (typically performed by multiple agencies) could be greatly simplified if trade details were available on a block within a blockchain. All involved parties could potentially read the details and perform their parts in the process without having to go through daily exercises, end-of-day processing, and numerous reconciliations.

• **Collateral management**. Reducing credit risk in financial transactions can require securing them with collateral – often through a complex set of steps that could be streamlined using a blockchain.

• **Payments processing**. Payment hubs are now looking into how to process those payments in a cheaper and more reliable fashion, and blockchain technology offers a potential solution.

• **Trade finance**. Typically, there are many parties and multiple reconciliation steps involved in trade finance processes, which can also require substantial transactional resources. Blockchain technology could make these processes simpler and less expensive.

• **Same-day settlement**. Companies are clearly moving toward reducing the number of days needed to settle a trade. With a blockchain, the full transaction can be completed within about 15 minutes from the time the initial message about the trade is sent.

• **Security servicing and processing**. With a blockchain, anybody who is subscribing to information about a security (for example, news about corporate actions or dividends) can immediately receive that information directly from the source.

• **Transfer agency and registrar functions**. Companies want to know who their shareholders are. Currently, when shareholder ownership changes, there is a long process before the company hears about it. With a blockchain, all of the steps in the process can occur on the same day.

• **Equilibrium and decentralization of players**. With a blockchain infrastructure, small players can play a larger role in markets that might otherwise be controlled by larger players (large custodians, depositories, transfer agencies and so on). Players will not need to invest a lot of money into the infrastructure, as everyone will use the shared infrastructure of the blockchain.

• **Private blockchains**. Anyone launching any of the use cases described here will need to decide whether they are comfortable using the same network used by today's bitcoins. The other option would be to build their own blockchain and allow only valid participants to use it, using a registration model instead of free entrance.

Let's take a close look at a specific example of how a blockchain could simplify a multi-step transaction process.

Creating a Blockchain-Friendly Infrastructure

Blockchain technology holds enormous potential for streamlining transaction processing. Making the most of this potential will require participants to develop an appropriate infrastructure.

With blockchains, there is no central data repository. Instead, the parties involved must collect the blocks of transaction information that are sent out to all subscribers. These blocks constitute an ever-growing amount of data (in fact, for bitcoins, there is just one blockchain of all transactions since the beginning, constantly increasing in size). Large-scale participants in this infrastructure will want to hold onto this data for their records – which means dealing with large amounts of data arriving in real-time.

Financial services firms not only need to store the incoming blocks, they also need to translate and validate the information. They need to maintain tables of security identifiers, validation information, cross-references, and so on. They need real-time transaction processing and analytical capabilities in order to leverage the advantage of immediate access to blockchain data for multiple transactions.

Any application that involves large-scale storing and processing of blockchain data will require a high-performance, scalable architecture. Some applications, such as bitcoin mining, require especially strong analytical capabilities and processing power. A blockchain-friendly architecture should also include streaming support, enterprise-level security features, and the high level of transactional consistency needed for financial systems.

Fortunately, in-memory computing is available with these features and is well-suited to this type of high-speed, big-data use case. Let's look at why in-memory computing makes sense for blockchain applications.

Blockchain and In-Memory Computing: Benefits and Use Cases

As noted earlier, blockchains can only grow, never shrink, since their integrity depends on no block being deleted. Large-scale blockchain subscribers must be prepared to quickly process, store, and analyze fast-growing amounts of transaction information arriving in real time. Because in-memory

computing involves keeping data in RAM for extremely fast access, with no disk-related slowdowns, it is faster than any other storage-based computing method.

For applications that require heavy analytics and real-time transaction processing of hundreds or even millions of transactions per second, the market is now moving from disk to in-memory computing. The reasons for this trend involve both performance and Return On Investment (ROI).

1000x Faster. The move from disk to memory is a key factor in improving performance. However, simply moving to memory is not sufficient to guarantee the extremely high memory-processing speeds needed at the enterprise level. Enterprise-level speed requires cluster computing, with multiple machines performing analyses at the same time, and parallel distribution of data. These capabilities are important for providing high availability, disaster recovery, and concurrency across systems – and they are all provided in the GridGain in-memory computing platform. Clients who have implemented the GridGain have found that they can process transactions 1000 times faster.

10x ROI Improvement. The cost of memory has dropped roughly 30% per year since the 1960s. Memory has become much more affordable in recent years. While it may still be slightly more expensive than disk, the performance is so much better that it improves ROI significantly. Clients who have implemented the GridGain in-memory computing platform have seen a tenfold or more improvement in their ROI.

Customer Case Study: Sberbank. One of the most noteworthy GridGain Systems financial services customers is Sberbank, the largest bank in Russia and the third largest in Europe. Sberbank was faced with a similar problem to the one currently facing companies who are transitioning from legacy systems and methods to real-time, synchronized transaction processing using blockchains. The bank was switching from a more traditional, brick-and-mortar setup – one in which people would come into their offices and manually process a limited number of financial transactions each day, during a limited time period – to a new world with online and mobile customers transacting with them 24/7.

The company forecasted future throughput requirements and determined that it needed to move to a next-generation data-processing platform to handle the expected transaction volume. Sberbank analyzed more than ten potential solutions from vendors in the in-memory computing space and found that the GridGain in-memory computing platform was the most comprehensive solution. The bank concluded that GridGain would provide the next-generation platform with a significant improvement in performance and scalability.

The GridGain in-memory computing platform provided several other important capabilities that Sberbank's next-generation platform would require such as machine-learning and analytics, flexible pricing, artificial intelligence, ease of deployment, hardware independence of cluster components, and a rigorous level of transactional consistency. Of particular importance was the ability to conduct integrity checking and rollback on financial transactions. Sberbank could not find that level of consistency with other in-memory computing solutions.

In a January 2016 article in RBC, Herman Gref, the CEO of Sberbank, said that the bank selected the GridGain Systems technology to build "a platform that will enable the bank to introduce new products

within hours, not weeks." He went on to state that the GridGain in-memory computing platform enables Sberbank to provide "unlimited performance and very high reliability" while being "much cheaper" than the technology used previously. Sberbank is using GridGain's in-memory computing platform to implement capabilities that could not be provided by the other vendors evaluated – a group that included Oracle®, IBM® and others.

GridGain Systems: A Leader in In-Memory Computing

As blockchain interest grows, many financial services firms are looking toward the GridGain in-memory computing platform. Of particular interest are GridGain's in-memory Data Grid, Compute Grid, Distributed SQL, Persistent Store, and streaming functionality. These features provide the performance, scale, and interoperability required for successful blockchain projects.

In addition to working with Sberbank on its next-generation system, GridGain is working with other banks that are now in a proof-of-concept stage or are close to moving blockchain applications into production. In addition, technology firms are evaluating GridGain as a means of offering Software-as-a-Service platforms around their blockchain implementations. GridGain's comprehensive platform contains a complete feature set that surpasses the capabilities of in-memory database point solutions, making it well suited to these types of financial use cases.

As a complete in-memory computing platform, GridGain helps users consolidate onto a single high performance and highly scalable big-data solution for transactions and analytics, resulting in lowered TCO. Advanced SQL functionality and API-based support for common programming languages enable rapid deployment. These features, along with the rapidly decreasing cost of memory, boost ROI for in-memory computing initiatives, enabling financial services companies to build less expensive systems that perform thousands of times better.

Leveraging the Power of Blockchains

Blockchain technology offers a compelling path forward for financial firms seeking to move away from a batch-processing and reconciliation-based model and toward real-time, synchronized transaction processing. Companies are embracing this digital-ledger technology bitcoin and non-bitcoin transactions. These companies require the performance and scale provided by in-memory computing to store, aggregate, and process large volumes of blockchain data at real-time speeds.

The GridGain in-memory computing platform overcomes the challenges of Bitcoin and blockchain projects with features that improve performance, scale, and interoperability. Data Grid improves the speed and scale of blockchain transactions. Compute Grid provides parallel processing for calculations required for building and updating blockchains. Compute Grid's parallel processing capabilities also speed up reconciliation and compliance checking. Distributed SQL allows companies to connect the GridGain In-Memory Computing Platform to other systems for analytics and reporting.

Combining an open source framework with enterprise-level features, the GridGain in-memory computing platform offers a high performance, scalable, and comprehensive solution. The GridGain in-

memory computing platform is an elegant and efficient solution to leverage the power of blockchains for streamlined transaction processing.

Chapter 3

Boost Performance of Financial Services Internet of Things (IoT) Projects with In-Memory Computing

Beyond all the buzzwords, the Internet of Things (IoT) is shaping up to be very big business. Imagine 26 billion devices—not counting computers and smartphones—creating data, communicating with the Internet, and resulting in $1.9 trillion in global economic value-add. Those numbers, reported by Gartner, echo a report from IDC which calculated $737 billion in spending on IoT equipment and services in 2016. Add to that an often-cited EMC prediction that by 2020, the digital universe will contain 44 trillion gigabytes of data, and the data management challenges of the future become readily apparent.

The financial services industry is already embracing IoT devices, collecting data that needs to be analyzed in real time and stored for historical analysis. Car insurance providers use telematics to monitor driver behavior, while health insurance providers are eager to connect with fitness tracking devices. Banks are installing location-aware ATMs that pre-load account information from approaching customers who are identified by their cell phones. Could an investment manager get a health alert from a client's wearable monitor and readjust the client's portfolio to limit risk as an emergency measure? All sorts of ideas are moving from the realm of the possible to the probable.

The platform for all this data collection, storage, and analysis must have several winning characteristics:

- Highly efficient sensors and devices
- Ubiquitous high-bandwidth network connectivity
- High availability
- Fast and scalable back-end storage, computational, and analytical systems
- Streaming data collection in near real time
- The best possible security
- The ability to adjust to variable workloads
- Interoperability

As IoT devices become part of our daily lives and generate huge amounts of streaming data, businesses are going to have to find a way to cope with the constant influx of data, and financial services companies will be on the leading edge. It is time for them to look at in-memory computing solutions.

In-Memory Computing: The Right Choice Now

In-memory technology has been around for decades, ever since developers realized that working with data stored in memory was faster than working with data stored on disk. Today, that speed means everything. Digital transformation has pushed companies closer to customers who demand real-time interactions, and the amounts of Internet traffic, connected devices, and data—all of which require instant analysis—continue to grow.

Luckily, the cost of memory continues to tumble—down 30 percent annually over time—so moving to in-memory platforms has gotten more cost-effective with each passing year. As a result, Gartner has

reported, the in-memory technology market will grow to $10 billion by the end of 2019, representing 22 percent compound annual growth over time.

The combination of these two trends means that the time is right for the financial services industry to leverage the improved performance and scale provided by in-memory computing to turn their IoT plans into action.

GridGain and Apache® Ignite™ Tackle Huge Data Volumes

The growth of the in-memory computing market includes in-memory databases, in-memory analytics, in-memory data grids, in-memory app servers, in-memory messaging, event stream processing, and analytic servers. The GridGain Systems and Apache Ignite in-memory computing platforms provide complete support for all of these components.

The GridGain and Apache Ignite in-memory computing platforms run everything in memory to yield the performance and scalability gains anyone managing huge IoT data streams will need. The GridGain Enterprise Edition difference is that it adds expected enterprise features such as better security, better network segmentation protection, data center replication, and rolling updates that allow updating of multiple nodes on a cluster without taking the whole cluster down.

FinTech Developers Can Take Advantage of In-Memory Benefits

GridGain has a wide range of customers in financial services industries including banking, insurance, and investment management, and in a 2016 survey, GridGain asked them what they were considering using in-memory computing for. The top uses: database caching, application scaling, high-speed transactions, and real-time streaming, all of which would serve an IoT project well. GridGain asked them where they would run such a system, and the two leading responses were on premise and in a private cloud, with AWS, Azure, and in another public cloud lagging. Luckily, the flexibility of GridGain and Ignite let them run anywhere.

Which of the following languages do the respondents use to manage their data? Java and SQL were the two leaders by far. C++, MapReduce, Scala, and .NET followed. It is interesting to see C++ rank this high. It is definitely a language that has earned favor in the financial services community.

Which data stores did they say they were likely to use as their persistent layer underneath GridGain and Apache Ignite? Oracle, MySQL, Microsoft SQL Server, and HDFS were the leaders.

Survey respondents also said that Data Grid, Compute Grid, and the In-Memory File System were the three most compelling features of the GridGain system, a response that makes sense given the importance of data scaling and the need for distributed computing platforms in financial services industries.

Digging deeper into GridGain's in-memory computing platform, we learn how those components work together. The powerful starting point is the I**n-Memory Data Grid**, inserted between the application and

database layers to cache the disk-based data from the RDBMS, NoSQL, or Hadoop databases in RAM, and deliver a performance increase up to 1,000 times. This in-memory key value store can replicate and partition data caches across multiple nodes and deliver elastic on-demand scalability to add new nodes. Distributed in-memory transactions can also be ACID-compliant. The Data Grid offers support for all popular RDBMS, with read-through and write-through and support for write behind. Setup remains completely flexible to address any unique use case.

When streaming data floods in from IoT sources, users can store it in Data Grid, and work with it with machine learning and all kinds of real-time analytics using the distributed computing architecture of Compute Grid. In-Memory Streaming and Complex Event Processing comes into play as users manage what is an unending data stream, setting event windows based on time or a specific volume of data and running operations or analytics even as the data continues to flow.

With all these features and tools enabled and running optimally, financial services developers can vastly improve their applications by:

- Processing high-speed transactions faster than ever
- Speeding up search dramatically
- Deploying better systems for the management of IoT components

Deploying In-Memory Computing for Optimal IoT Implementations

Sberbank, the largest bank in Russia and Eastern Europe and the third largest in Europe—it has 16,000 branches across 11 time zones—was expecting significant growth in its transaction volume, from 30 transactions per second a decade ago to up to 4,000 today. It wanted to eliminate its vertical software stacks for things like payments and mortgages and minimize its legacy dependency on Oracle by migrating to an open-source data grid architecture for its next-generation platform. The bank needed to introduce new products in hours, not weeks, and the platform had to have virtually unlimited performance and very high reliability while lowering costs and reducing human interaction.

The GridGain solution Sberbank implemented offers a very high level of redundancy and high availability and has the security levels Russian regulators demand. In testing, the bank was able to manage up to one billion transactions per second on an array of ten Dell R610 blades with one terabyte of memory, assembled at a cost of just $25,000.

On the IoT side, GridGain worked with a company specializing in Internet-connected devices to address a large real-time data processing challenge. The company's connected energy meters deliver data from millions of homes, and it forecasted that its existing database infrastructure would soon no longer be able to keep up. Not only did it need to ingest massive amounts of data constantly, it also had to overlay current electricity rates and perform compute-heavy calculations on the fly in order to produce regular analytical reports combining current state and recent history.

This company was committed to keeping as much of its database infrastructure as it could, so the solution was to add GridGain on top of it, creating a fast space in which to do the required real-time data processing and analytics and providing as much scalability as the company needs—from 100 to

1,000 nodes if necessary—as its business grows. Today the company can store a week of its current data entirely in RAM, vastly improving its ability to deal with the never-ending rush of data it must process moment by moment.

IoT Will Mean Money in the Bank

In-memory computing gives financial services companies the speed, scalability, and flexibility they need to build IoT-based applications and services that will serve their current clients in new and profitable ways while attracting new clients as well. No matter what kinds of Internet of Things implementations come to fruition in the future, it is clear that the high-performance experience and real-time analytics across a variety of platforms that in-memory computing offers will be the firm foundation on which they are built.

Chapter 4

Banking on In-Memory Computing with GridGain and Apache Ignite

Addressing the challenges highlighted in this eBook requires speed, scalability, availability, security and flexibility. In short, they need distributed in-memory computing. In-memory computing eliminates the disk access bottleneck that slows down applications built on disk-based databases. An in-memory computing platform enables users to process transactions 1,000 times faster than disk-based solutions and enables scale out to terabytes of in-memory data by adding new nodes to the cluster.

In-memory technology has been around for decades. Until recently, however, the cost of RAM made in-memory computing practical only for the highest value applications. The cost of memory continues to fall, dropping an average of 30 percent per year, which makes in-memory computing platforms economical for a wider range of use cases each year. Gartner projects that the in-memory technology market will grow to $10 billion by the end of 2019, a 22 percent compound annual growth rate.

In-memory computing platforms include key features that are now essential for many financial applications. These features, available in Apache Ignite and the **GridGain in-memory computing platform**, include an in-memory data grids with strong distributed SQL capabilities, in-memory database capabilities, streaming analytics, and native integrations with a variety of other open source projects including Apache® Kafka™, Spark™, Cassandra™ and Hadoop™.

GridGain and Apache Ignite are deployed as an in-memory computing layer between the application and data layers. The products work with any RDBMS, NoSQL or Hadoop database. The **In-Memory Data Grid** with strong distributed SQL capabilities is a key-value store which can replicate and partition data caches across multiple nodes and deliver elastic on-demand scalability. Distributed in-memory ACID transactions are also supported. The Data Grid offers support for all popular RDBMSs, with read-through and write-through and support for write behind. Setup is flexible to address unique use cases.

The **In-Memory Database** leverages the systems distributed SQL capabilities. It is horizontally scalable, fault tolerant and ANSI-99 SQL compliant. It also supports all SQL, DDL and DML commands including SELECT, UPDATE, INSERT, MERGE and DELETE queries and CREATE and DROP table. The SQL syntax is ANSI SQL-99 compliant. GridGain can use any SQL function, aggregation, or grouping. GridGain supports distributed SQL joins and allows for cross-cache joins. Joins between partitioned and replicated caches work without limitations while joins between partitioned data sets require that the keys are collocated. GridGain supports the concept of fields queries as well to help minimize network and serialization overhead. The in-memory distributed SQL capabilities allow users to interact with the GridGain platform not only with the usage of natively developed APIs for Java, .NET and C++ but also using standard SQL commands through the GridGain JDBC or ODBC APIs. This provides a true cross-platform connectivity from languages such as PHP, Ruby and more.

The **In-Memory Compute Grid** enables distributed parallel processing of resource-intensive compute tasks. It offers adaptive load balancing, automatic fault tolerance, linear scalability, and custom scheduling. Built around a pluggable SPI design, it offers a direct API for Fork-Join and MapReduce processing.

The **Distributed SQL** is horizontally-scalable, fault-tolerant, and ANSI SQL-99 compliant with support for all SQL, DML and DDL commands such as SELECT, UPDATE, INSERT, MERGE, and DELETE queries or CREATE or DROP tables. It is a mature, in-memory solution to supplement or replace a disk-based RDBMS. Geospatial support is built into the product and all the communication to the SQL grid is done through ODBC and JDBC APIs without custom coding.

The optional **Persistent Store** feature in the memory-centric Apache Ignite architecture is a distributed disk store that transparently integrates with GridGain as an optional disk layer. It may be deployed on spinning disks, solid state drives (SSDs), Flash, 3D XPoint or other similar storage technologies. Persistent Store allows organizations to maximize their return on investment by establishing the optimal tradeoff between infrastructure costs and application performance by adjusting the amount of data that is kept in-memory.

The **In-Memory Service Grid** provides control over services deployed on each cluster node and guarantees continuous availability of all deployed services in case of node failures. It can automatically deploy services on node startup, deploy multiple instances of a service, and terminate any deployed service. It is a load-balanced and fault-tolerant way of running and managing services across the grid.

Stream analytics establish windows for processing and run either one-time or continuous queries against these windows. The event workflow is customizable and often used for real-time analytics. Data can be indexed as it is being streamed to make it possible to run extremely fast distributed SQL queries against the streaming data.

In-memory Hadoop acceleration provides easy-to-use extensions to disk-based HDFS and traditional MapReduce, delivering up to 10 times faster performance. GridGain and/or Ignite can be layered on top of an existing disk-based HDFS and used as a caching layer offering read-through and write-through while the GridGain Compute Grid can run in-memory MapReduce.

Many enterprise financial services firms have adopted GridGain's in-memory computing solution to improve performance and functionality. Clients including Apollo Global Management, Barclays, Cambridge Associates, Citi, Fidelity, and ING have all taken advantage of GridGain's distributed computation and in-memory computing speed.

Conclusion

With the tight regulatory environment, competition from traditional and non-traditional industries, customer demands, and cost pressures that companies are facing today, e-commerce initiatives require big data technologies that make processes and transactions much faster and more efficient. Large companies accumulating massive amounts of data need to be able to perform analytics on that data in real time in a cost-conscious manner to ensure a good user experience. Many are finding in-memory computing platforms such as GridGain and Apache Ignite to be a key strategy for meeting these challenges.

Additional Resources

For more information about topics covered in this eBook, the following resources are available from the GridGain website:

eBook: In-Memory Computing for Financial Services: Part One

Datasheet: The GridGain In-Memory Computing Platform

Case Study: Misys Uses GridGain to Enable High Performance, Real-Time Data Processing

Article: **How Russia's Oldest Bank Found Itself on the Edge of In-Memory Computing in CIO Magazine**

Webinar Recording: Modernize and Accelerate Payment Solutions with In-Memory Computing

Webinar Recording: Enable Bitcoin and Blockchain Technology with In-Memory Computing

Contact GridGain Systems

To learn more about how GridGain can help your business, please email our sales team at sales@gridgain.com, call us at +1 (650) 241-2281 (US) or +44 (0) 7775 835 770 (Europe), or complete our contact form to have us contact you.

About GridGain Systems

GridGain Systems is revolutionizing real-time data access and processing by offering an in-memory computing platform built on Apache® Ignite™. GridGain solutions are used by global enterprises in financial, software, e-commerce, retail, online business services, healthcare, telecom and other major sectors, with a client list that includes Barclays, ING, Sberbank, Misys, IHS Markit, Workday, and Huawei. GridGain delivers unprecedented speed and massive scalability to both legacy and greenfield applications. Deployed on a distributed cluster of commodity servers, GridGain software can reside between the application and data layers (RDBMS, NoSQL and Apache® Hadoop®), requiring no rip-and-replace of the existing databases, or it can be deployed as an in-memory transactional SQL database. GridGain is the most comprehensive in-memory computing platform for high-volume ACID transactions, real-time analytics, web-scale applications and HTAP. For more information, visit gridgain.com.