MEDICI™

Why Financial Services Should Look to NoSQL

Executive Summary

The digitization of financial services has been an ongoing process for the past few decades. The transition has not always been smooth and easy from legacy mainframe systems of the 60s and 70s to today's cloud-based financial services. Customers are increasingly adopting digital channels to access banking services and other offerings. **Neobanks have launched that are completely digital, without having any branch banking.**

Financial services can no longer offer generic solutions that target a wide range of consumers. Be it an individual working in a multinational corporation or a student, a neighborhood business, or a large-scale enterprise catering to a global populace, the needs of each client are unique. To address this diversity, institutions are consuming data from multiple sources inside and outside the enterprise. These include structured data from sources like enterprise systems, market systems, and government databases, as well as unstructured data from social networks and media.

The emphatic acceptance of modern platforms and **technologies such as NoSQL databases (which are powering tech giants like Facebook, Amazon, and Google) has provided the opportunity to tackle these challenges effectively**. This report explores the various facets of NoSQL databases and how their growth has played a crucial role in transforming several financial services business areas. FinTech companies such as neobanks, online trading platforms, and payment platforms are already incorporating NoSQL databases in their architectures and are reaping the benefits of being agile and resilient.

This report also investigates the several advantages that **NoSQL offers over traditional databases for a financial institution**. The ability to handle structured and unstructured data formats, ease of connecting legacy systems with new, speedier front-end systems, horizontal scaling, and faster implementation are some of the points that are explored in detail. It also analyzes how each of these advantages translates to better performance in the various functions of a financial institution ranging from fraud prevention, better customer service and customer 360, targeted marketing, data analytics, and real-time digital identity verification.

Financial institutions are increasingly looking to move operational and delivery models from physical to digital platforms as more and more customers prefer these channels. **NoSQL will play a critical role in enabling this shift and in ensuring the high quality of services** that their new, digital clientele expects from them.

MEDICI™ Copying or distribution without written permission is prohibited

NoSQL Databases in the Changing World of Financial Services

Copying or distribution without written permission is prohibited

Digital transformation in financial services and the changes imposed by it

Financial services institutions, especially banks, have long been struggling with legacy systems. When implemented in the ‘60s and ‘70s, they ensured that these financial institutions were ahead of the curve in terms of computerization; however, now they have become massive IT burdens that are dragging these behemoths into the ground. **One of the major issues is the reduced pool of IT staff who are well-versed in these legacy systems**.

Issues faced in changing these legacy systems:

COST

Most of the time, firms have to go for a complete overhaul, and that involves massive costs in terms of both hardware and software procurement.

CHANGE MANAGEMENT

Most of the top management in financial services firms are still resistant to change. They fear short-term problems over long-term benefits and are unwilling to take the call to bring about this change.

BREAK-IN OPERATIONS

Every second is crucial for the financial services industry. The industry is apprehensive that any wholesale change to the systems could result in a long-term break-in operation and potential loss of business.

SKILLED MANPOWER

There is a reduced pool of IT staff who are well-versed in these legacy systems. Their absence makes it difficult to engage in IT transformations effectively.

Other Issues with Banking Legacy Systems

- Big banking systems implemented in the ‘70s were written primarily in COBOL, which was introduced in the ‘60s.
- In the ‘70s and ‘80s, banks were the leaders of tech innovation and saw the birth of products such as ATMs, BACS, and international card payments.
- With the arrival of the internet, customers wanted access to financial transactions from the comfort of their homes, and banks responded by investing in client-side software.
- Gradually, they worked out that client-side software was not the right solution compared to the standards established for internet traffic and figured out that a website worked way better for a fraction of the cost.
- The Big Data Revolution that has taken hold in recent years is now causing financial institutions around the world to enhance their capacity to access and mine data from all sorts of sources.

Copying or distribution without written permission is prohibited

The emphasis of digital transformation in financial services during the last 20 years has focused on the following points:

1. The emergence of digital delivery of financial services products and services
2. Increased automation of the financial services enterprise (internal digitalization focused on process automation for efficiency and cost optimization)
3. Increased reliance on data for analytics-based business decision-making
4. The emergence of cloud —both platforms and technologies—brought in the need to change legacy architectures to take advantage of the agility and lower cost of cloud computing.

Many banks are still struggling with shifting from the legacy systems and are severely behind the curve. Neobanks are hot on the heels of traditional banks, and with more digital avenues available, customers will be frequenting physical branches even less. This is the time for banks to go all out in technology adoption and ensure that they usher in a new era in innovation and customer satisfaction.

What is a NoSQL database?

One of the changes seen in IT systems of financial institutions is the adoption of NoSQL databases. The next few sections will examine what a NoSQL database is and how its adoption can help in the digital transformation of the financial services industry.

NoSQL is usually referred to as not only SQL, non-SQL, or non-relational databases. NoSQL databases such as key-value pairs are built for prime throughput versus relational databases with relative dependence. Using loose dependencies and quick indexes **NoSQL databases are perfect for streaming analytics and IoT applications** because data can quickly be stored and referenced from multiple, disparate data sources.

NoSQL databases represent an approach to data management and database design that's useful for very large sets of distributed data. They encompass a good range of technologies and architectures, seeking to resolve the scalability and Big Data performance issues that relational databases were not designed to deal with. **NoSQL is particularly useful when an enterprise has to access and analyze massive amounts of unstructured data or data stored remotely on multiple virtual servers within the cloud**.

Copying or distribution without written permission is prohibited

Types of NoSQL databases

Columnar Database

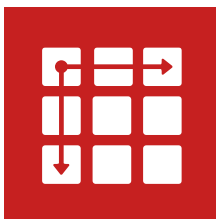

In column-oriented NoSQL databases, data is stored in cells grouped in columns of data instead of rows of data. Columns are grouped into column families that contain a virtually unlimited number of columns that can be created at runtime or while defining the schema. In these databases, read and write functions are done using columns rather than rows. The benefits of storing data in columns over relational DBMS are fast search, access, and data aggregation.

Column databases are primarily used in systems that maintain counters, content management systems, blogging platforms, services that have expiring usage and systems that require heavy write requests.

Graph-based NoSQL Database

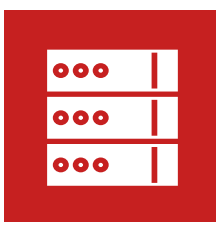

Graph databases are essentially created upon the Entity-Attribute-Value model. Entities are also known as nodes, which have properties. It is a very agile way to describe how data relates to other data. Nodes store data about each entity in the database. Relationships describe a relationship between nodes, and the property is the node on the opposite end of the relationship.

Graph-based NoSQL databases are well-suited for graph-based search, network and IT operations, and social networks.

Key-value store NoSQL Database

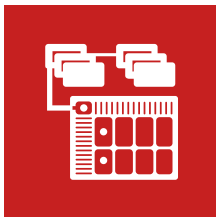

Key-value stores are the simplest NoSQL data stores to use from the perspective of APIs. The client can either get the value for the key, assign a value for a key, or delete a key from the data store. The value is a hold that the data store just stores without caring or knowing what is inside. It is the responsibility of the application to understand what was stored. Since key-value stores always use primary-key access, they generally have better performance and can be easily scaled.

However, new-age Big Data applications have sparse datasets, and traditional databases cannot handle them as effectively. In the case of facilitating efficient selection of a subset of the dataset based on business constraints, **hybrid NoSQL databases use secondary indexes by providing alternative access paths to the base records**. Alternate to a scan, these special index structures can help identify the records that qualify, and then only those records are retrieved from the base table. Secondary indices are defined over one or more attributes and are often constructed over non-primary-key attributes.

Key-value databases are primarily used for storing user session data, maintaining schema-less user profiles, and user preferences. A lot of data generated in financial institutions is time-series data as the industry is highly regulated, and adequate audit trails are required to be maintained. NoSQL databases are used to handle such data that has a timestamp and includes a time order.

Document Store NoSQL Database

Document store NoSQL databases are similar to key-value databases in that there is a key, and there is a value. Data is stored as a value, and its associated key is the unique identifier for that value. The difference here is that, in a document database, the value contains structured or semi-structured data. This structured/semi-structured value is referred to as a document and can be in XML, JSON, or BSON format.

Document store databases are well-suited for e-commerce platforms, analytics platforms, content management systems, blogging platforms, and more.

MEDICI™ Copying or distribution without written permission is prohibited

Growth of NoSQL databases

NoSQL technology was originally created and used by Internet leaders such as Facebook, Google, Amazon, and more, which required database management systems that could read and write data anywhere in the world while scaling and delivering performance across large data sets and billions of users.

However, a large number of institutions worldwide are now adopting NoSQL databases to handle bigger datasets more efficiently and to ensure analytics and real-time decision-making using the same. **Industries such as media, financial services, e-commerce, research, and analytics are increasingly moving towards NoSQL databases** to handle their data management because of its major advantages, such as lower cost, distributed computing, and high scalability.

Some of the major NoSQL database providers include:

AEROSPIKE | amazon web services | DATASTAX | Microsoft | Couchbase

Google | MarkLogic | mongoDB | neo technology | Objectivity

As per a report from Allied Market Research,

- The global NoSQL market is expected to reach **$22.08 billion by 2026.**
- The industry is expected to grow at a **CAGR of 31.4%** (2019-26) [1].

As more and more people start transacting on the internet, the volume of data generated increases manifold, and companies are looking for data management solutions that can handle them better.

Along with efficiency, NoSQL databases also offer cost-effective modes to handle data as compared to traditional databases. **A Forrester report had concluded that NoSQL products saved companies more than 50% of the cost compared to conventional DBMS** [2]. NoSQL database implementation is relatively easier and typically uses less expensive servers to manage rising data and transactions. In contrast, RDBMS databases are costlier and use larger servers and storage systems. So the storing and processing data cost per gigabyte in the case of NoSQL is several times less than the cost of RDBMS. These advantages are causing enterprises to switch over from relational databases. According to Gartner, by 2022, 50% of existing proprietary Relational Database Management Systems (RDBMS) instances will have been converted or be in the process of conversion [3].

MEDICI™ Copying or distribution without written permission is prohibited

Where NoSQL Databases Fit in Financial Services

NoSQL databases are better equipped to handle larger data sets

Some applications need simple object storage, whereas many others require highly complex and interrelated structure storage. NoSQL databases provide the required support for a range of data structures ranging from simple binary values to complex hierarchical structures.

NoSQL databases can handle smaller binary values, lists, maps, and strings at high speed in key-value stores. Slightly more complex, related information values can also be handled by key-value systems with complex data types or be grouped in column families within big table clones. A web of interrelated information can also be managed flexibly and related in triple and graph stores in a NoSQL database.

NoSQL databases are increasingly being used for Big Data analysis. This is because NoSQL databases typically follow the BASE approach instead of the ACID approach of relational databases (although this is now beginning to change). By being Basically Available (BA), it guarantees the availability of the data at all times. With Soft state (S), the state of the system could change over time, and Eventual (E) consistency ensures that the system will eventually become consistent once it stops receiving input. Legacy NoSQL databases, for the most part, give up the A, C, and D requirements, and in return, they improve scalability.

Ease of handling both structured and unstructured data formats

RDBMS cannot be considered the best solution for all situations as they cannot meet the increasing growth of unstructured data. Traditional databases cannot process unstructured and unpredictable information. However, modern applications demand an occurrence-oriented database. These have to be highly flexible and operate on a schema-less data model.

If a financial services institution only plans to pull data similar to an accounting spreadsheet, i.e., the basic tabular structured data, then the relational model of the database would be enough to fulfill the business requirements. However, current trends demand just the opposite: storing and processing unstructured and unpredictable information. This information comes from multiple sources. With the increasing size of the database or increasing number of users, RDBMS using SQL suffer from major performance and scalability issues, thereby making real-time unstructured data processing more difficult.

On the other hand, **NoSQL is a flexible and cloud-friendly approach to process unstructured data with ease dynamically**. With the availability of several mobile and web applications, it is very common for financial institutions to have millions of users who will generate a lot of unstructured data.

NoSQL databases simplify the representation of multi-level hierarchies and nesting using the JSON (JavaScript Object Notation) format. In the dynamic schema universe where changes happen every hour, it is impossible to stick to the 'get it right first' model, which was a success with the outdated static schema.

Web-based enterprises such as Amazon, Facebook, and even modern-day neobanks, require a database with NoSQL capabilities that can match up with changing data models, thereby offering them greater flexibility in operations while maintaining resiliency and uptime. They have the added advantage of not having to deal with legacy systems, unlike traditional banks. This enables them to work in real time and offer a unique customer experience not seen before in the financial world. From instant account opening to the quick transfer of money and personalized services, neobanks have caught the imagination of customers. A lot of this is enabled through real-time processing of data and business analytics. NoSQL databases can help traditional financial services organizations and neobanks alike to enhance their offerings further and innovate on products and services using real-time data analytics.

MEDICI™ Copying or distribution without written permission is prohibited

High-performance handling designed for real-time feedback at web-scale

The main purpose of keeping a database in-memory is performance. Reading and writing data that is in-memory is typically faster than data stored on a flash drive or a disk. As a result, the user does not have to wait for disk I/O in order to update or query data with two notable exceptions.

For in-memory systems, data is added to the main memory with the help of specialized indexing data structures. This data is always available in-memory but is also persisted to disk with logs and database snapshots. Lastly, the ability to read and write data so quickly and efficiently in an in-memory database enables mixed transactions and read/write workloads (although write workloads can be much slower and tend to bog down legacy in-memory NoSQL systems).

Many real-time applications in the modern world need the performance of NoSQL database structures. This helps ensure that decision-makers in the financial services industry have the most relevant and timely information when interacting with a customer, vendor, or the management in their own company. With this information in hand, they can prepare and present an accurate picture to the concerned parties rather than having to wait for information to be returned or basing their decision on outdated information. Having this data on hand so quickly provides a sure-shot competitive advantage for enterprises using tools such as customer relationship management solutions.

Built to connect legacy systems with newer and faster front-end systems

Many traditional financial institutions still use legacy systems for various functions in their organizations. Some of these systems have large amounts of data, and the process to convert them involves humongous effort and high costs. As a result, many of these organizations continue to retain these systems. However, they work with modern databases in areas such as front-end operations that have a lot of web-based interactions. **NoSQL offers a powerful way to monitor key performance indicators in real-time, which helps with decision-making, SLA adherence, and forecasting**.

Under a legacy approach to business intelligence, data from such online transactions are likely to languish in data warehouses. However, with a NoSQL database in place, that same information could instead be accessed for providing real-time insights into consumer behavior. For example, in a digital front-end for a traditional bank, key-value stores (for storing user sessions), graph stores (for tracking customer spend), document stores (for payment transactions), and column families (large scale web analytics and support for reports) can all be utilized through NoSQL.

Firms have to keep the following points in mind while linking legacy systems with NoSQL databases:

1. Identify integration points in the legacy stack, as integrations can happen at the Object Relational Mapping or database layer without any additional overhead

2. Ensure there is a strong architecture in place to enable the seamless integration of NoSQL and in-memory computation systems

Copying or distribution without written permission is prohibited

Built to drive two-paced development of modern architectures

The massive growth in mobile, Big Data, and cloud technologies has greatly changed market dynamics in every industry, including the financial services industry, driving the convergence of the digital and physical worlds, and changing customer behavior. A new approach is required to achieve such a scenario, **one that allows web-scale innovation in order that technology teams can meet changing business requirements** while enabling existing systems to continue running reliably, securely, and efficiently.

Can scale better horizontally as the data grows

NoSQL databases manage the sharding of a database across several servers. Hence, if data storage requirements increase, one can continue to add inexpensive servers and connect them to the database cluster, making them work as a single data service. This is known as horizontal scaling, and since it follows the denormalization concept where it can store duplicates, there is no need for a single point of failure in a NoSQL database. **Providing durability and high availability of a NoSQL database by using less expensive hardware and storage is one of NoSQL's major advantages.**

In a relational database, we would need to purchase more powerful and thus increasingly more expensive hardware to scale up vertically. If we were to increase the amount of data storage, we would have to increase the cost of the hardware needed as well substantially.

Easier and faster to implement compared to traditional databases

NoSQL databases have increased in popularity with the rise of Big Data-based applications. Unlike relational databases, **NoSQL databases are much cheaper to scale, capable of managing unstructured data, and better suited to the current agile development approaches**. The advantages of NoSQL technology are compelling, but replacing a legacy relational system can be daunting.

NoSQL databases are often open-source; users can download the software and build their applications on it for free (albeit with either scale or feature limitations and (or) lack of support). They are also inexpensive and easily available across multiple software vendors. **NoSQL database implementation is easy and typically uses cost-effective servers to manage exploding data and transactions. In contrast**, RDBMS databases are more expensive and use big servers and storage systems. Hence the storing and processing data cost per terabyte in the case of NoSQL can be many times less than the cost of an RDBMS.

Support for better performance handling

The performance of a relational database model can be enhanced with more powerful hardware. However, as discussed earlier, this process is costly and time-consuming. The continuous upgrade process of hardware is the problem that NoSQL databases are trying to solve. Instead of having to purchase new servers with more memory, **NoSQL databases apply a scale-out model where we can easily add new servers to the cluster**. This enables the organization to grow without having to replace their existing hardware investment. This benefit alone makes it worth considering NoSQL for many modern use cases.

In the question of multi-row transactional support, it has been observed that many use cases today do not require updating multiple rows and tables in a transactionally guaranteed manner, i.e., as ACID transactions, which is why NoSQL works so well, in comparison. Also, since NoSQL databases use a denormalized model in which all information about an entity is presented in a single record, there is no requirement to update multiple rows.

Copying or distribution without written permission is prohibited

Opportunities and Business Benefits That Can Be Derived from NoSQL Databases

Copying or distribution without written permission is prohibited

Improved fraud prevention and financial crime monitoring

Fraud prevention in financial services is a great example of an application that benefits from real-time analytics. **It is a very crucial aspect for banks, trading institutions, and payment processors**. Failure in these areas results in losses: both financial and reputational, and both false positives and false negatives can have massive business implications.

The global fraud detection and prevention market size stood at **$19.90 billion in 2019** [4] **and is projected to reach $110.04 billion by 2026** [5]. This represents a recognition that fraud is becoming incredibly challenging and can only be addressed with cutting-edge technology.

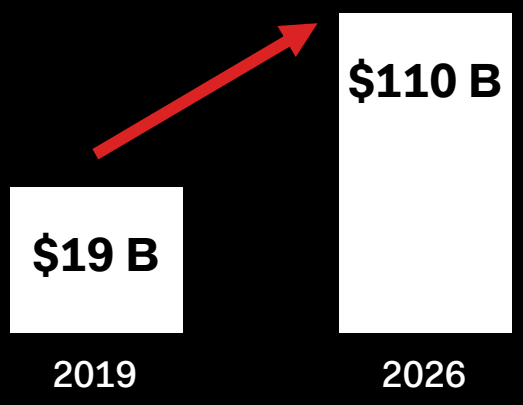

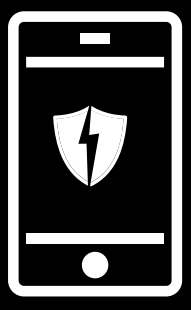

RSA reports a **600% increase in mobile fraud over the last three years, till 2018** [6]. As more transactions occur on mobile channels, the element of fraud occurring on the platform is only going to increase.

These days, fraud detection is handled by:

- Analytics
- The ability to apply data to detection through anomaly identification
- Predictive classification
- Clustering

Fraud detection works on various decisions based on knowing the customer, monitoring transactions, and identifying patterns. Many companies are now building transactional analytics systems for fraud to complement their existing architecture.

A NoSQL database offers real-time decision-making on large datasets. Latency plays a crucial role here, and the lower the latency, the more the tasks that can be performed to check the various parameters involved in fraud detection. **NoSQL databases enable real-time analytics on live transactions by removing the latency associated with moving data from operational databases to data warehouses for analytical processing**.

Copying or distribution without written permission is prohibited

Real-time digital ID validation

ID verification has gained increased significance in the global, digital world. As more and more of the global population connects to the web for their various requirements, it is imperative that firms identify genuine customers against anti-social elements.

In 2021, **cybercrime damages could reach $6 trillion** [7]. The Global Risks Perception Survey by the World Economic Forum identified cyber-related issues, such as cyberattacks and data fraud or theft, within the list of top 10 long-term risks globally.

ID verification firms across the world connect with multiple digital ID databases to ensure that clients are able to verify the identity of the people they interact with, instantly. Many of these clients belong to the financial services industry, who have to ensure that ID fraud does not happen at various levels of interactions with customers and clients. The ID verification firms have to instantly make API calls to these databases and display the results in almost real-time. NoSQL databases reduce read latency in such situations and ensure that these organizations are able to handle real-time customer trust decisions virtually, thereby eliminating false positives and massively enhancing fraud detection.

Personalized offerings using improved real-time data analytics

Personalization is the buzzword in the financial services industry these days. **Customers do not want generic products and services that cater to a wider population**. Each individual's financial requirements are different, and they expect their bank or any other financial institution to offer them solutions matching their needs.

Customers are now asking for flexible credit card interest rates that vary based on the repayment and spend capacity. Insurance policies have to be tweaked based on localized risk factors, and investment professionals have to design bouquets based on individual risk appetite. Banks can use data from multiple sources to offer personalized loan products as well. For example, they can assess the risk of lending to an online vendor by analyzing their inflows from a payment gateway statements, reviews of their product on the online portal, and marketing reach based on the metrics of their social media handles.

For this, organizations must collect all possible information about their customers to prepare a comprehensive user profile of the individual. Those profiles will often have recent user behavior, segments loaded from an analytics system, partner cookies, social media engagements, online shopping preferences, and more. **Companies need robust data management tools such as NoSQL databases to store and access these multiple sources**. These will, in turn, enable the organization to make real-time analytics on user preferences and ensure that products and services are provided to cater to the specific needs of the customer.

Copying or distribution without written permission is prohibited

Targeted marketing

Targeted marketing is closely tied to the personalization aspect mentioned above. Publishers, advertising operations, and agencies use various platforms and tools for buying, selling, delivering, and targeting display advertisements. These video and mobile ads need to occur in real-time as the customer can arrive and leave a digital property in a fraction of second.

Targeting marketing has become the most effective medium for brand engagement from social media sites, displaying ads based on viewer preferences to recommendation engines on streaming platforms and e-commerce sites. This is only possible through large-scale data crunching offered by NoSQL databases using data sourced from user profiles, cookies, and social media engagements. Managing massive amounts of audience segmentation data exchanges behind this transaction between players such as buyers, sellers, and ad exchanges requires low latency and efficient data management on the part of databases.

Pre-trade assistance

As stated, **NoSQL databases can be used to offer personalized services by studying customer behavior through data from multiple sources**. Another use case of the same is data analytics for trade decisions and risk assessment. Pre-trade assistance uses analytics to study information in real time from various sources, both highly structured market data, and unstructured news information and research. Some of these data sources are also geospatial, such as commodities data.

Instead of just displaying this data, **real-time analytics can be performed using NoSQL databases that provide a holistic view of understanding a particular instrument**. Traders can use this data to arrive at an objective analysis of the instrument's performance and make informed decisions accordingly. Traders can also be alerted about significant changes derived from this information, thereby giving them an edge over others, as they are relying on holistic analytics rather than unstructured data.

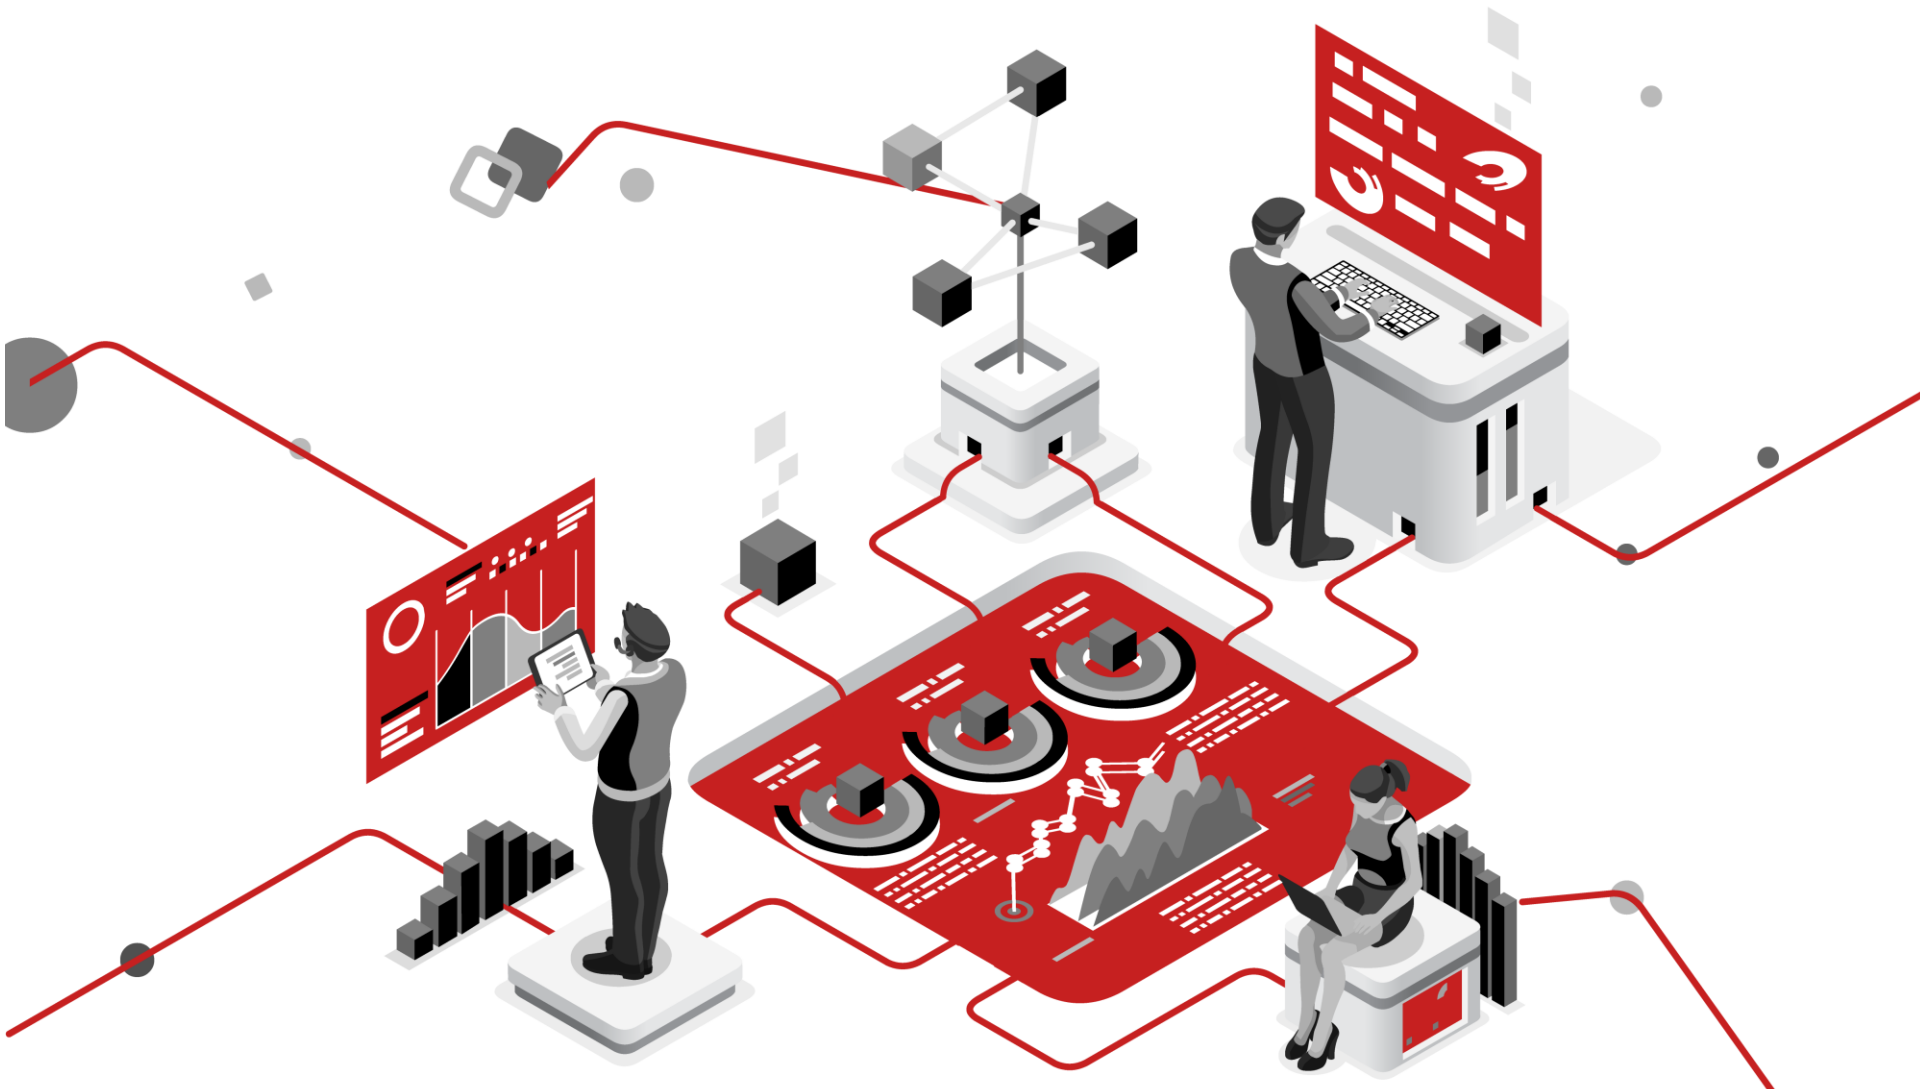

 Copying or distribution without written permission is prohibited

Conclusion

Conclusion

Technology is going to play a massive part in the new environment for Financial Services. Customers who are accustomed to instant updates through social media will expect a similar service from their financial providers.

NoSQL databases bring immense value when an enterprise has to access and analyze massive amounts of unstructured data or data that is stored remotely on virtual servers within the cloud. NoSQL databases will be one of the pioneering technologies for financial institutions in the coming years, enabling them to succeed in the digital environment.

As firms around the world look at cost-effective measures to run businesses, technology will be the main focus. With the growth in data, it will become uneconomical for firms to spend on scaling databases vertically by adding more compute and storage capacity.

Instead, with **NoSQL, applications can store their data in purpose-built databases, which is especially true in the case of unstructured data**. This will also aid in robust data analytics for improvements in business performance and customer service.

Modern, cutting-edge platforms such as **NoSQL databases will help accelerate innovation to achieve its full potential so that the industry meets its consumers' demands and expectations**.

MEDICI™ Copying or distribution without written permission is prohibited

Bibliography

1. NoSQL Market Report, Allied Market Research
2. The Forrester Wave: Big Data NoSQL, Forrester
3. Magic Quadrant for Operational Database Management Systems, Gartner Research
4. Fraud Detection & Prevention Market Size, Share & Trends Analysis Report, Grand View Research
5. Fortune Business Insights Report
6. Quarterly Fraud Report, RSA
7. Cybersecurity Ventures Official Annual Cybercrime Report

MEDICI™ Copying or distribution without written permission is prohibited

About MEDICI

Since 2013, MEDICI has been pioneering the definition and organization of the new global FinTech industry for the benefit of financial institutions, startups, and investors.

Over the years, it has been a story of numerous Firsts in FinTech!

We were the first independent source of data-driven research dedicated to covering FinTech innovation globally, published every day since August 25th, 2013, now offering an archive of 5000+ insights for the benefit of this new industry.

We built the industry's first content curation and ecosystem collaboration platform in 2015. Today, this proprietary technology supports the world's largest FinTech community of 200,000+ globally across 1000+ enterprises, 13,000+ startups across 65+ sub-segments covering every hub of innovation in every continent.

Website: www.goMEDICI.com | **Twitter:** @gomedici

Thank You

Aditya Khurjekar
CEO & Founder
ak@goMEDICI.com
@khurjekar

Amit Goel
CSO & Founder
amit@goMEDICI.com
@amitTwitr

Global Contacts

Salil Ravindran
Head, Digital Banking
salil@gomedici.com

Giuseppe Marchese
Head of Business Development, Europe
giuseppe@gomedici.com

Subhagini Chaudhary
Head of Sales, Southeast Asia
subhagini@gomedici.com