Whitepaper

February 2021

Data-Driven Telecom

Rethinking database technology platforms for a distributed future.

Authors: Francis Haysom, Robert Curran

Working with

Appledore RESEARCH

Couchbase

Published by Appledore Research LLC • 44 Summer Street Dover, NH. 03820

Tel: +1 603 969 2125 • Email: francis.haysom@appledorerg.com• www.appledoreresearch.com

Figures and projections contained in this report are based on publicly available information only and are produced by the Research Division of Appledore Research LLC independently of any client-specific work within Appledore Research LLC. The opinions expressed are those of the stated authors only.

Appledore Research LLC recognizes that many terms appearing in this report are proprietary; all such trademarks are acknowledged, and every effort has been made to indicate them by the normal USA publishing standards. However, the presence of a term, in whatever form, does not affect its legal status as a trademark.

Appledore Research LLC maintains that all reasonable care and skill have been used in the compilation of this publication. However, Appledore Research LLC shall not be under any liability for loss or damage (including consequential loss) whatsoever or howsoever arising because of the use of this publication by the customer, his servants, agents or any third party.

Cover photo by Marcelo Moreira from Pexels

Contents

INTRODUCTION **4**

FROM PROCESS-DRIVEN TO DATA-DRIVEN **5**

DATA-DRIVEN TELECOM **6**

DATA-DRIVEN GROWTH **9**

THE LEGACY SYSTEMS CHALLENGE **10**

BECOMING DATA-DRIVEN **11**

Scalability 11
Resilience 11
Distributability 11
Analytics 12
Performance 12
Applicability 12
Compatibility 12
Support 13

CONCLUSION **13**

INTRODUCTION

In the last 20 years, the telecom industry has experienced greater and faster change than at any time in its history. At the start of 2021, that pace shows no sign of slowing down.

Communications service providers (CSPs) have seen previously secure revenue streams eroded and been forced to rethink basic assumptions about what customers want, and where they are willing to go to get it. Major vendors have also experienced a downturn from the heady days of the 1980s, with wafer-thin margins, less predictable revenues, and yet significant ongoing needs to invest in R&D.

Developers of key business and operations support software (BSS/OSS), together with their systems integration partners, have been engaged for 10 or more years in programs of transformation – intended to enable a step-change in the productivity, agility, and profitability of CSPs. Yet few people would argue that this effort is anything close to complete, despite some successes.

And the increasing use of software to provide critical network functions, from identity management to network policies and security, continues to see new specialist vendors innovating in the marketplace.

In recent years, three primary agents have accelerated the pace of change: mass-market broadband, cloud computing and the smartphone. These three have completely transformed how customers define their experience – and their expectations.

Yet these are not what will determine the forward trajectory for telecom.

While each continues to advance on their own terms with ever-greater speed, reach and capabilities, their combination creates a completely new operating model for telecom. This new operating model creates new demands on how telecom relies on – and what it looks for – in its most critical IT platforms.

And at the heart of this new operating model is **data**. More specifically, the ability to deliver the right data, at the right time, and in the right place.

This new context is highlighting the technical (and commercial) limitations of database platforms selected during an earlier, simpler generation in telecom. For everything that telecom now aspires to be – dynamic, distributed and disaggregated – a new generation of database technology is now needed.

FROM PROCESS-DRIVEN TO DATA-DRIVEN

Data has been an integral part of the development and operation of telecom networks. But in the last few years, major factors have changed:

- The **volume and variety of data** of all kinds has massively expanded – not only content, but also telemetry and activity data from network devices and end-user devices.
- **Automation** has become a driving factor across the telecom operations landscape.
- Recognition of the **need to integrate isolated silos** of customer, network and operational data.
- Greater **flexibility** to choose exactly where and how data is used within telecom's own IT architectures.
- The need for greater **speed and agility** in developing new, hyper-personalised data-centric applications
- The move to **cloud** infrastructure and web-scale architectures

Part of the response has been improvement in sophistication of database technologies, upgrading relational database technologies but increasing investments in emergent NoSQL databases. (It is worth remembering that the "no" in NoSQL stands for "not only".) These extend the capabilities of conventional RBDMS with better ways to store, retrieve, update and manage large amounts of data, especially where that data not always well-defined or well-behaved. "Big Data" technologies provided a solution to some use cases but tended to focus largely on applications within a homogeneous data set – pattern recognition, trends, visualisations and dashboards. Operational systems, however, require platforms that are dynamic, distributed and scalable.

Telecom's response to the automation imperative was initially led by projects to mechanise, and to re-implement manual processes using automated workflow systems. But not fundamentally to change underlying data or data management technology. More recently, data (rather than just process) has become more integral to automation efforts, and the scope has broadened from mechanising specific business processes. Targets for automation now spans a wide range of user and customer-centric telecom processes across front office, back office and networks, such as:

- Customer complaint routing and handling processes
- End-user device failure diagnosis and resolution processes
- Selection and presentation of ads to content subscribers based on viewing behavior
- Proactive replacement of network equipment based on historical trends or detection of network patterns
- Security, authentication of users, devices or equipment

But the rise of cloud computing forces telecom to see data and its associated functions in a new light. Because in addition to the challenges of size and scale, and the need to automate, now the

location of data becomes a variable factor in how fundamental processes and functions work, and services are made possible. The degree to which data can now be distributed represents a new dimension for creating business value, reducing hard costs (such as software licenses), and improving customer experience.

The result is that in telecom, the technology for managing data can no longer be considered a discrete, localised, technical, project-by-project decision. The capabilities (or limitations) of technologies for managing data have – and should have – a direct bearing on critical KPIs across telecom, not only in the IT function.

The application of specific database technologies, the performance achieved, or the economies of scale envisaged, are no longer only for IT consideration, but have a direct impact on customer experiences, live service KPIs, network operating costs, profitability and even brand reputation.

This is what is means to be **data-driven**. To recognise that without the capabilities required for managing what data now means in this genuinely new generation in telecom, the full potential of this new era simply will not be achieved.

DATA-DRIVEN TELECOM

Thanks to cloud computing, the data that telecom relies on can be (or in some cases may need to be), processed at different locations along a compute spectrum, from data centre to one or more cloud locations to network edge or on a device.

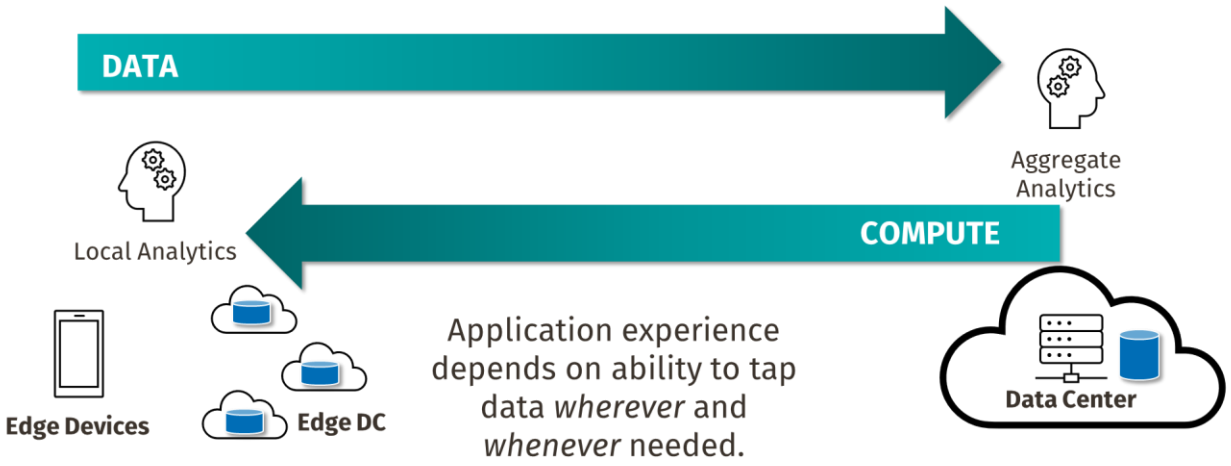

Figure 1: From Centralized to Distributed Data

New insights can be gained by backhauling data from the edge of the network, aggregating it in a central location for mass analysis, and distributing the results (as data or algorithm updates) back into the network or to devices on its edge.

This is what creates the generational change in the understanding of data. The CAP Theorem provides a framework for considering which of three priorities to compromise on in a distributed system:

- Consistency (data is always accurate, characterised through support of ACID transactions when required)
- Availability (the application will always receive an answer, tolerant of network and internet availability)
- Partitioning tolerance (the application can operate even in the event of node failure(s))

To some extent, telecom hedged the partitioning factor by created a distributed network, but as the network increasingly becomes a dynamic, virtual entity and not a fixed one, the hedge is no longer valid.

Modern telecom's uses of data now represent a variety of priorities and compromises, according to different application needs – not only data, but also latency and responsiveness. As such, the flexibility of the underlying database to adapt to different needs is an important factor.

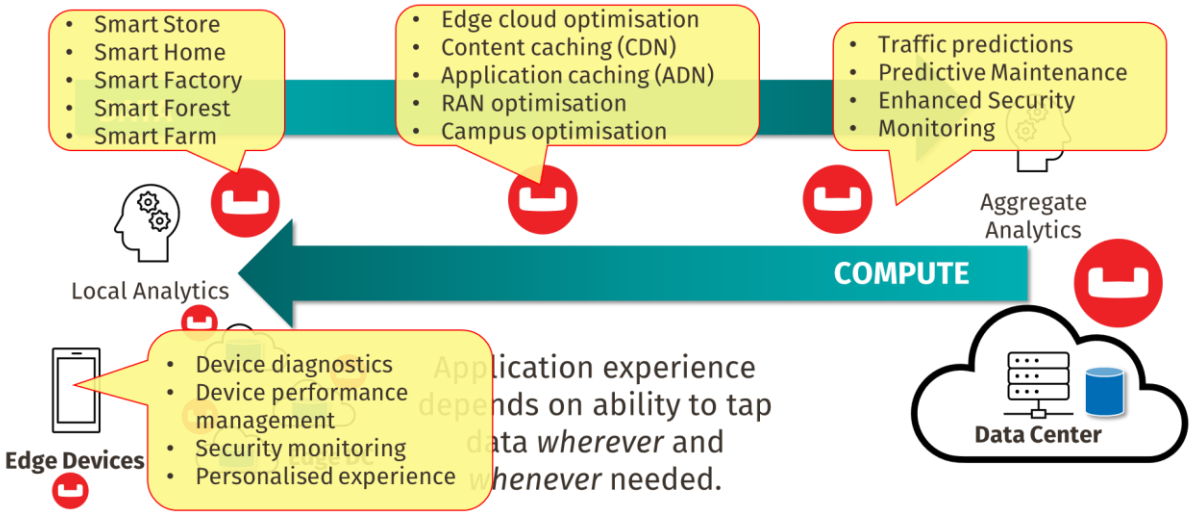

Figure 2: Example applications leveraging distributed data & compute for hyper-personalization

The generational change required can be seen by dividing telecom's data into distinct domains:

Domain	Data Examples	Example Application
Customer Experience	• Subscriber data – behaviour, patterns, commercial information, contact history	• Self-service handset upgrade • Media consumption personalization • Enhanced customer profiling based on contact history
Services & Applications	• Service data (in-service data; performance, experience) – esp video quality-related • Application data – in-app data points: speed/frequency of interaction, user feedback/ratings (direct/indirect), user attentiveness…	• "Smart Factory", with video monitoring of production line quality fed back to control speed • "Smart Forest", gathering local environmental data to provide early-warning of fires
Network	• Network data – health check, device intelligence/prediction, class of device, vendor information, maintenance (physical), maintenance (digital) • Network performance measurements – latency, jitter, quality • Field service data	• Content caching and bookmarking • Radio Access Network optimisation • Edge network optimisation • Traffic hotspot predictions • Predictive maintenance • Dynamic network control • Power optimisation
Devices	• Device data – status, performance, temperature, location, speed…	• On-device diagnostic applications
OSS & BSS	• Network inventory, fulfilment, and assurance • Billing, rating, and mediation data	• Renewal and TCO reduction • Move to cloud • Remote field service inventory • Alarm store & analytics

Telecom has always struggled to integrate data between these domains. Yet now, the ability to do that will unlock additional intelligence, revenue potential and cost savings. For example:

- Using subscriber trends (BSS) to change individual RAN parameters to reduce power consumption (Network)
- Using device behavior data (Device) to predict new product uptake (BSS) or application suggestions
- Using network KPI data (Network) to alter the availability of service features (OSS and BSS)
- Using KPI data from neutral edge hosts to multiple clouds to optimise caching strategy
- Maintaining user service among small edge networks (for example, a car) when internet or networks are unavailable

These are new kinds of applications, calling on data sourced from anywhere in the network, processed in dynamically determined locations in the cloud, in time frames that vary from the ultra-

low-latency of 5G to the 2G speeds that continue to suffice for millions of connected meters, detectors and probes.

DATA-DRIVEN GROWTH

A recent report on the major growth opportunities for telecom identified a dozen areas of major new opportunity for operators and vendors. In most, a generational change in how data is managed is necessary.

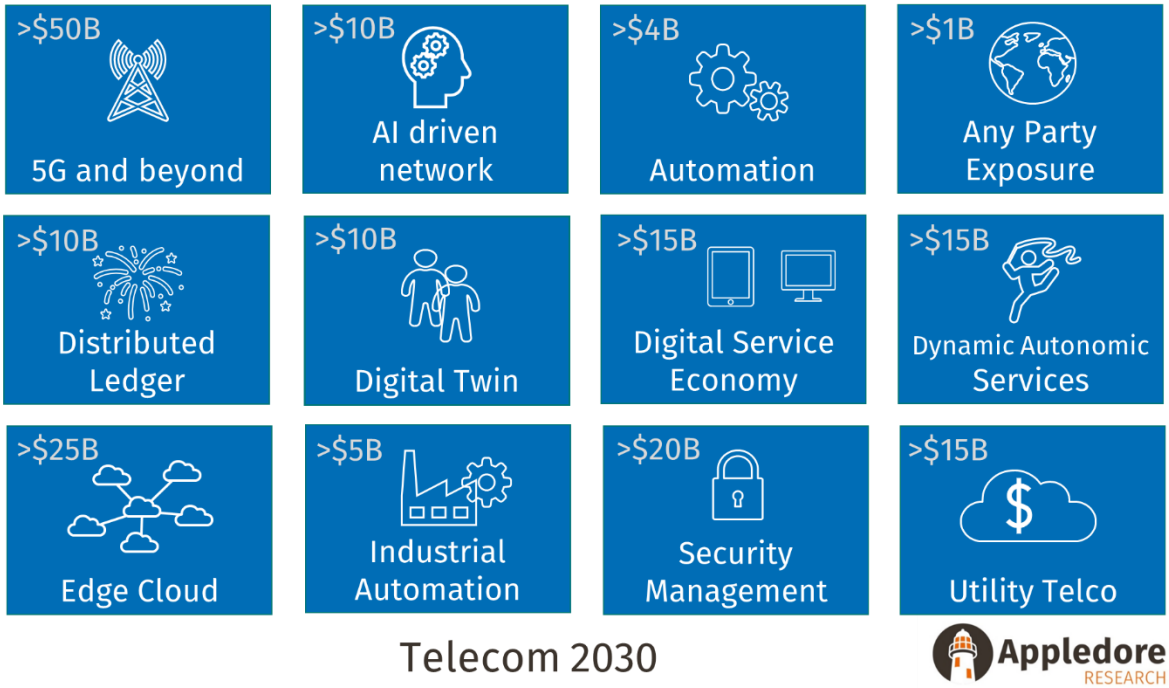

From an IT perspective, 5G completes the work that 4G started, in providing an architecture for mobile networks which is fundamentally centred on data, software, and applications. Dynamic, data-driven networks are only possible if the data that drives them can be relied on.

A data-driven network must combine both real-time and historical data from network devices in order to make intelligent choices and inferences about how to reconfigure the network and/or devices.

"Automation" is now on telecom's agenda as never before. The ability to make instantaneous change to networks means that telecom cannot afford processes that depend on human intervention. There simply are not enough people to do the job as frequently as it needs done. Assessing the condition of thousands of servers, switches, routers; deciding which telemetry data matters when, at what granularity level?

And Automation depends on data. More specifically, on having the right data in the right place at the right time to act on. And not only data, but the related analytics, security, management, performance functions that fuel automated decision-making.

The commercial potential of distributed compute resources to a micro-edge compute cloud demands both application and core network data to be gathered to make intelligent decisions about how and where to balance application load.

The use of a telecom network "digital twin" further illustrates how data is used in new ways with applications operating on data to constantly adapt the behavior of a device, subsystem or whole network, in response to constantly updating data:

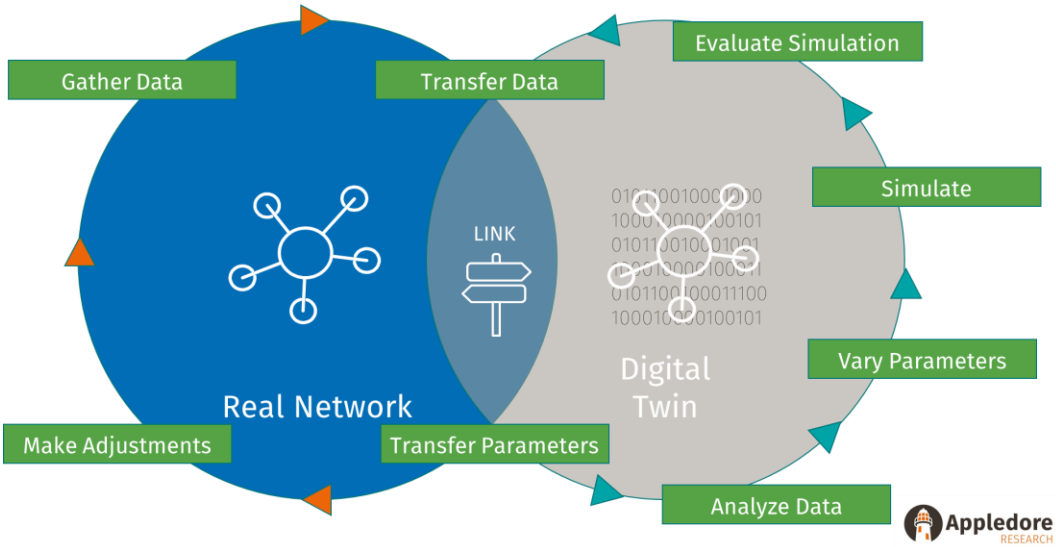

Figure 3: "Digital Twin" data cycle

In the digital twin, information gathered from across the network in real time is aggregated, analysed and used as the basis of additional prediction, modelling and simulation – in order to feed change instructions back to the live network. The availability of data around this cycle is critical to making the digital twin work.

THE LEGACY SYSTEMS CHALLENGE

Billing and Operations Support Systems (OSS) are sometimes regarded as the "crown jewels" of a telecom operator. These are the systems that maintain the most essential data for a telco. And that data is constantly being added to, with every order, call made, app downloaded, support chat, text message, video watched and GPS calculation.

The convergence of IT and telecom has established a radical departure from conventional telecom practices and related systems. These core areas of operations have always been characterised by large data volumes (digitising an entire network with every cable, connector, piece of card, rack, shelf, node and server equipment), and the rate at which that data changed was in line with processes to make physical changes in the real world.

With virtualisation of telecom, changes to the network can and must be made hundreds of times more quickly, and – critically – driven by data. The availability and reliability of data, close to where decisions are processed is a critical requirement of telecom's most exciting fields, such as edge computing, 5G, private networks and for the application of analytics within a wide range of industry verticals that telecom supports.

The underlying database platform technology choices made a decade or more ago could not have envisaged the radical changes that have taken place since. However, enlightened operators and ISVs are transforming their applications and architectures to meet the new challenges of distributed and dynamic networks, massive and complex data, dispersed and decentralized device networks, with a primary focus on user experience and operational automation as their core differentiators.

While this is a significant technical challenge, modern database platforms offer flexibility in how to approach such a transformation. This can initially be by acting as an overarching cache to satisfy the high concurrency demands placed on existing systems, or in construction of new "smart" applications in support of modern edge computing capabilities of new 5G architectures. This can then be extended as legacy platforms are de-supported, refactored or become uneconomical to retain unless they are moved to the cloud.

BECOMING DATA-DRIVEN

With all this in mind, how does this generational change alter the criteria for considering database platforms? We believe the following characteristics are the most important:

Scalability

In telecom, the inevitable trajectory for data of all kinds is growth. In earlier decades, the response to greater data was to upgrade servers and software licenses. That is no longer either economical nor practicable. Technical and commercial models that scale up, and do so gracefully, are essential to avoid painful, disruptive and expensive changes later.

Clearly, "as-a-Service" models offer more flexibility – but these must be backed by a modern underlying architecture that delivers true scalability in a distributed up, out and across multi-cloud environments.

Resilience

Instantaneous access to data is now integral in the provision of live services in telecom, an industry which has been defined by its insistence (and delivery) of "five-nines" reliability: uptime of 99.999%. When that service is real-time communication between emergency services, or between robots transporting hazardous materials in a factory, the resilience of data provision to network failures is critical.

Distributability

The trio of "Analytics + Mobility + Broadband" is often cited as a transformational combination in telecom. But it underplays an important underlying capability: the ability to locate data and compute resources at the most appropriate point along a continuum.

More compute resources are being located closer to the edge of the network (and end-user devices themselves have become immensely more powerful.) The rationale for this is largely to improve application experience – specifically to reduce delay. There can be other benefits in terms of data residency, costs of network capacity.

4G and 5G networks specifically encourage flexibility (or at least, an agnostic approach) to where compute resources live. This allows applications to be distributed across a combination of on-device and cloud locations. But the data could also be distributed – dynamically – to suit the needs of that application.

A customer service call could involve reviewing years of historical transactions on demand, within an ad hoc query. A request for mobile user's profile could happen anywhere within a global network. A 5G network slice could be created to provide real-time analysis of a video feed of a disaster to help locate survivors.

The ability to distribute data without creating single points of failure is an important feature of the move to data-driven telecom.

Analytics

As highlighted already, large amounts of data can now be accumulated in many source locations. Rather than backhaul data to a central location, analysing it in situ will reduce delay and deliver a better experience. As such, database platforms with built-in analytics capabilities are highly desirable.

Performance

Raw performance of database platforms is already a key requirement. Now, that extends beyond read/write cycles to the ability to sync, replicate and rebalance data between nodes in a cloud-based or other distributed network. Also, the ability to performance-match the services of the data platform to both the requirements of the application and to its underlying infrastructure is a key to optimizing for both performance and cost.

Applicability

As we have illustrated, telecom offers a huge variety of ways to realize value from application of new data, and new applications for existing accumulated data. Database platforms that can be applied to a range of use cases – from high-availability core network functions to latency-sensitive edge applications – will allow operators, vendors and integrators to service a wide range of needs with a common skill set. This will reduce application development time and cost. And when focused on the user experience of customers and consumers, can provide a positive impact to revenue.

Telecom operators, ISVs and SIs should resist the temptation to make expedient tactical or localised decisions based on a limited set of parameters, such as convention, apparent cost, or licensing structure. As this paper has highlighted, data and its applications are of far greater importance in the dynamic, cloud-based telecom future. Expedient choices today will constrain future flexibility and business potential.

Compatibility

Although data-driven telecom is a radical change, it will not be achieved overnight. New database platform technologies should offer a gradual introduction path, from acting as a caching layer on top of legacy databases, to eventual replacement in order to gain performance and flexibility benefits.

Support

As IT and telecom have converged, so open-source programs have gained popularity. However, at the same time, there is an increasing awareness of the need for clear chain of accountability and resolution path on sensitive or mission-critical network components. Telecom's organisational culture has long been defined by its five-nines imperative, and that will not change. Operators in particular should consider the relative technical merits of open-source database platforms against the needs to manage business risk.

CONCLUSION

Telecom has awakened to the importance of being a data-driven industry. In a previous generation, data was a by-product of a business process; a way to record past events. As such, managing it was largely an archiving function requiring technologies that needed only be good at storing large amounts of information, retrievable by humans.

But by 2021, it is clear that data, not just connectivity, is the basis of creating new value in telecom: great experiences, new insights, smarter change. The explosion in available data, and change in its significance, should prompt a re-assessment of the underlying technologies used not only to store it, but to keep it secure, to analyse it, to send it where it needs to be – with flexibility, at speed and scale.

The agility and flexibility that the telecom industry desires will be enabled, or limited, by the capabilities of database platforms to deliver data quickly, securely and reliably.

Telecom vendors and operators should regard data, especially data that drives personalization for consumers and customers. as fundamental to their ability to execute, differentiate and succeed. Consequently, they should take a strategic perspective on the choice of technologies for a future that is inherently uncertain, but which offers powerful and valuable new possibilities – driven by data: the right data, delivered to the right place, at the right time.

Working with

E: info@appledorerg.com

www.appledoreresearch.com

44 Summer Street Dover, NH. 03820, USA