White Paper

Reducing risk and increasing resilience in RTGS payment systems

Contents

Summary

Why RTGS matters

Real time gross settlement (RTGS) reduces counterparty credit risk in payments by settling transactions gross, one by one in real time, instead of netting payments between counterparties and settling the net amount either at the end of the business day, or at regular intervals throughout the business day. RTGS is provided, primarily for high value payments (HVPs), by RTGS systems operated mainly by central banks. Payments are still netted by automated clearing houses (ACHs), primarily for low value payments (LVP), but even ACHs settle net payments in the RTGS system. Through services provided to custodian banks by central securities depositories (CSDs), the securities markets also rely on RTGS systems to provide final, irrevocable settlement in central bank money of the cash leg of securities transactions. Because they are vital to cash and securities settlements, and also play a vital role in the implementation of monetary policy, RTGS systems are systemically important.

Regulatory pressure for resilient RTGS systems

Because RTGS systems are systemically important, the central banks which operate them must ensure that they are resilient enough to withstand a variety of threats to their security and integrity. These include natural disasters, loss of essential services, data corruption, cyber-attacks, unavailability of staff, component malfunction, terrorism and war. The 24 principles for financial market infrastructures published in April 2012 by the Committee on Payment and Settlement Systems (CPSS) and the Technical Committee of the International Organisation of Securities Commissions (IOSCO) emphasise not only final settlement in central bank money, in real time, as the global standard, but also the need for operational contingency plans that guarantee continuity of service through both catastrophic and marginal disruptions.

Best practices in RTGS resiliency planning

Uninterrupted provision of RTGS services requires a high degree of resilience. In major markets, every RTGS system is supported by a complete back-up site, mostly run on so-called "hot-standby mode" enabling it to capture transaction information continuously, and so take over functions immediately in the event the primary site is disabled. A second site is nevertheless vulnerable, particularly if it operates on the same technology as the primary site.

RTGS Resiliency Approach

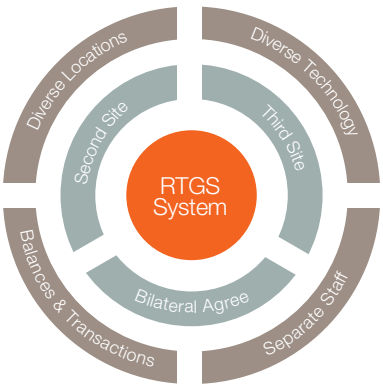

For this reason, some RTGS operators have built a third site, but this is an expensive option because it delivers additional resilience only if it is operated by separate staff on different technology in a remote location. At present, an alternative in the event of the loss of the secondary as well as the primary site is reversion to bilateral settlement between counterparties. This entails a mix of manual and automated processing, restricting the volume of payments that can be processed, and reintroducing a degree of credit risk. There is also no means of capturing the balance of payments made and pending at the point of failure of the primary and secondary sites, leading to disputes between counterparties.

The characteristics of a truly effective resiliency plan therefore include affordability; a rapid cut-over to the new service; a geographically remote facility; reduced reliance on local staff; technical diversity; independent data storage; sufficient capacity to support existing volumes of business; minimal impact on users; the availability of the service throughout the period of disruption; and, most importantly of all, the ability to capture a clear view of the intra-day balances at the point of failure, or to recreate it rapidly once the primary and secondary sites have failed.

The solution

For the last five years, SWIFT has worked with a group of seven central banks to design a shared RTGS system back-up service which meets best practices, including the tests of affordability, capacity, rapid implementation, minimal impact on users, geographical remoteness and technical diversity. Called the Market Infrastructure Resiliency Service (MIRS), it makes use of SWIFT technical platforms, storage facilities and messaging formats to capture transaction balances continuously, and so guarantee the ability for the operator to open the MIRS back-up service to the RTGS participants within no more than 2½ hours by providing a clear view of the settlement position at the point of failure of the primary and secondary sites. The service is easy to use and can operate for as long as a disruption persists, whether this is a matter of days, weeks or months.

1. Why RTGS matters

1.1 The benefits of RTGS

Real Time Gross Settlement (RTGS) is a clumsy term for a crucial process in the financial markets. This is the reduction of counterparty credit risk by the delivery of cash or the delivery of securities in exchange for cash, instantaneously and without the netting of the obligations outstanding between the parties.

Since the 1980s, the central banks which operate payment market infrastructures (PMIs)[1] around the world have gradually adopted RTGS for the settlement of high value payments (HVP). Their private sector equivalents which settle low value payments (LVP) are also gravitating towards RTGS.

In RTGS settlement, credit risk is reduced because cash is transferred between banks continuously in real time, transaction by transaction. Every payment is settled finally and irrevocably in central bank money, obviating the need to settle obligations between banks in batches on a net basis.

1.2 What is an RTGS?

The role of a PMI is to provide predictable and secure multilateral payment services to banks and their corporate and retail clients, usually within a single country, but sometimes across several countries within a region. They tend to divide into two broad groups. The first are HVP systems, which settle a relatively low volume of high value and high priority payments. The second are LVP systems, which are also known as Retail Payment Systems (RPS), because they net relatively high volumes of low value and low priority payments.

There is a further distinction to be made between HVP systems. Not all HVP systems settle on a gross basis in real time (RTGS). Some settle on a net basis, in which case they are technically described as High Value Payment Deferred Net Settlement (HVP DNS) systems. This is because settlement of transactions does not take place instantaneously but is instead deferred until transactions can be aggregated into batches, and the sums owed by one bank to another netted into a single net payment, made either at the end of the business day or at regular intervals throughout the business day. The net settlement typically takes place in central bank money at the RTGS.

LVP or RPS systems tend to net transactions in a fashion comparable with HVP DNS systems. Operated mainly by automated clearing houses (ACHs), they aggregate and net transactions between banks, and then settle net amounts between banks in central bank money at the RTGS either in a single payment at the end of the business day or in multiple payments made at regular intervals throughout the day.

Although a variety of net settlement systems persist, more than half the PMIs in the world are now RTGS, and even net settlement systems ultimately settle in RTGS (see Chart 1).

It follows that RTGS systems are crucial to the settlement of both HVP, LVP and CSD transactions. In fact, the purpose of every RTGS is to provide final, irrevocable settlement of transactions in a specific currency, usually through the transfer of the reserves held by banks at the central bank. They act on payment instructions, and settle transaction by simultaneously debiting the account of the paying bank and crediting the account of the receiving bank.

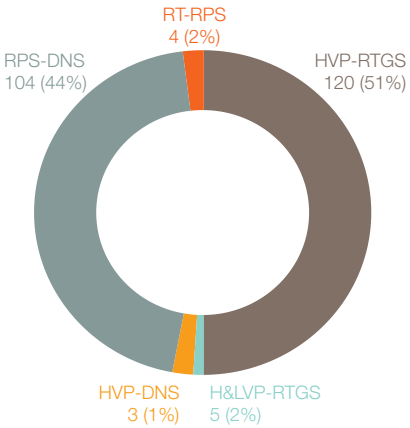

Chart 1: PMI systems in the world today

Reserves are a vital tool of monetary policy. They are the cash balances that banks are required to hold at central banks, both to limit the ability of banks to lend deposits without limit, and to guarantee the stability of the financial system by ensuring banks can always settle their obligations to each other. This makes RTGS an essential tool for every central bank in managing the stability of the financial system, because it is a means by which it can inject and withdraw liquidity (see Chart 2).

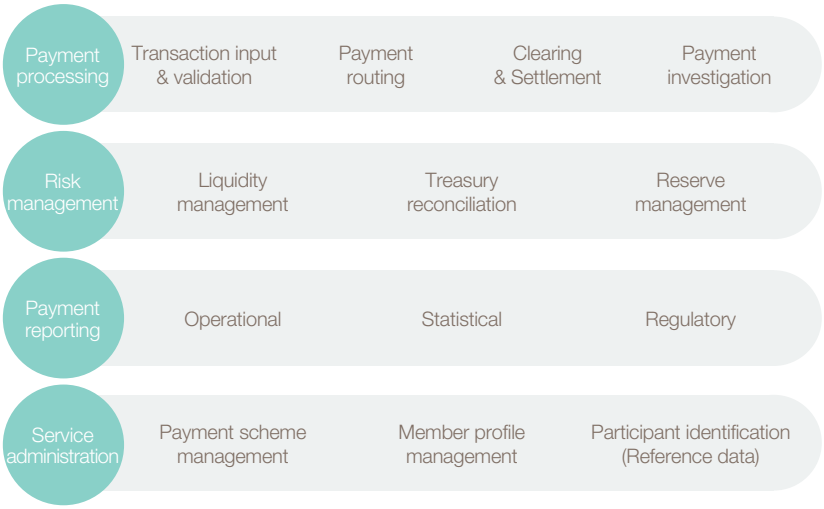

Chart 2: The key functions of an RTGS

1) Payment market infrastructures (PMI) is a Financial Market Infrastructures (FMI) - defined by the Committee on Payment and Settlement Systems (CPSS) and Technical Committee of the International Organisation of Securities Commissions (IOSCO), Principles for Financial Market Infrastructures, Bank for International Settlements, April 2012 - that operates a payment system, which is a set of instruments, procedures, and rules for the transfer of funds between or among participants; A payment system is generally categorised as either a retail payment system (RPS) or a high-value payment system (HVPS).

1.3 How RTGS reduces credit risk

RTGS is now regarded by central banks as essential to the reduction of counterparty credit risk. According to the International Bank for Reconstruction and Development (IBRD, or World Bank) the volume of payments settled by major RTGS systems doubled between 2006 and 2009, while the average value of RTGS payments increased by two fifths.

The official preference for settlement finality in central bank money was recently reinforced by the Committee on Payment and Settlement Systems (CPSS) and the Technical Committee of the International Organisation of Securities Commissions (IOSCO). In April 2012, they published 24 Principles for financial market infrastructures.[2] Principles 8 and 9 set final settlement in central bank money, in real time, as the global standard.

An HVP represents a significant credit risk if one party defaults. By guaranteeing final and irrevocable settlement at the central bank, RTGS mitigates this risk. By settling payments between multiple banks, it also eliminates the need for banks to settle transactions bilaterally. Banks are the principal users of RTGS, and most connect directly, though smaller banks often connect indirectly through a larger bank.

But the benefits of credit risk mitigation through RTGS are not restricted to banks. They extend to all participants in financial markets, whether they are active in the money or securities markets, and whether they are linked to the RTGS system directly or indirectly (see Chart 3).

Retail ACHs, which net large volumes of LVP before initiating a single high value net settlement payment in an RTGS system, have increased the frequency at which they settle from once to multiple times a day. In the most advanced cases, payments are taking place every 15 minutes or so, and the market trend is to increase the frequency still further.

The central securities depositories (CSDs) that deliver securities against payment use RTGS systems to settle the cash leg of transactions.

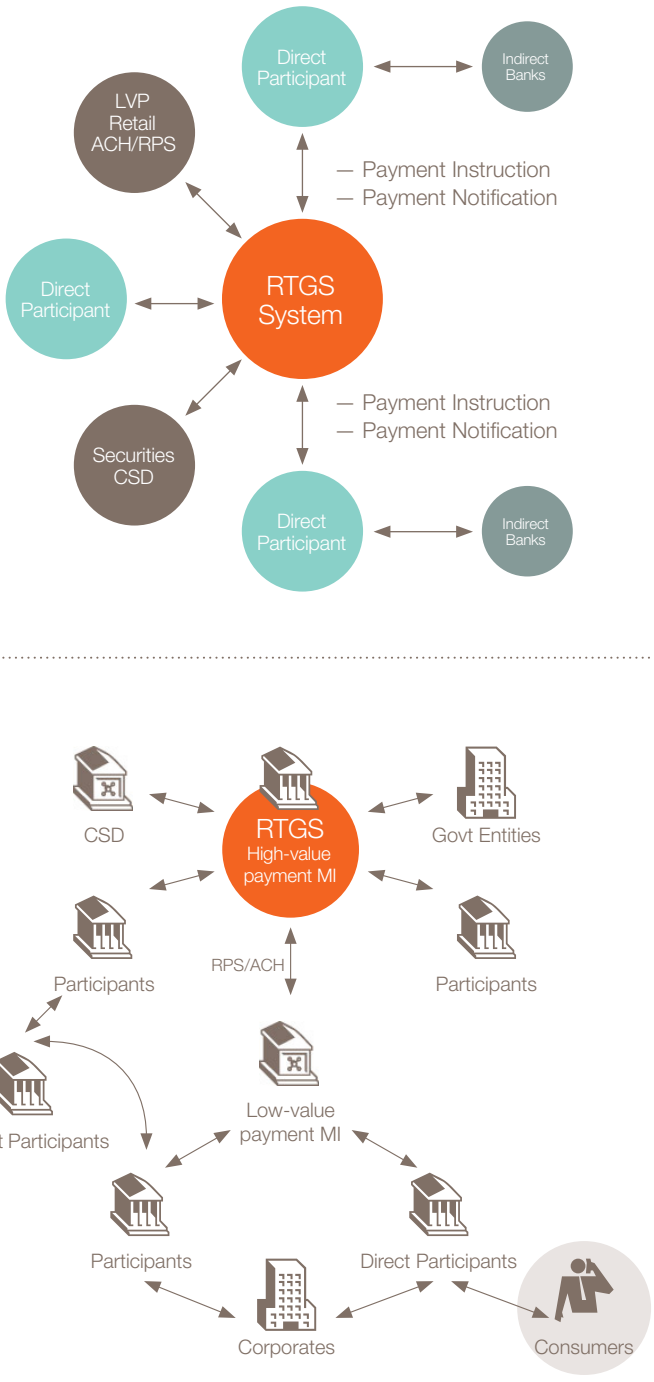

Chart 3: How RTGS benefits all market participants

2) Committee on Payment and Settlement Systems (CPSS) and Technical Committee of the International Organisation of Securities Commissions (IOSCO), Principles for Financial Market Infrastructures, Bank for International Settlements, April 2012.

They debit or credit the accounts at the central bank of custodian banks, which act on behalf of the broker-dealers and fund managers which buy and sell equity and fixed income securities in the markets. Cash payments driven by securities transactions are an increasingly important source of transactional activity in RTGS systems.

The connections between the securities markets and RTGS systems are further intensified by the fact that central bank money is available only against reserves or eligible collateral, such as cash or government securities. To clear and settle trades in central bank money, custodian banks must either use the cash balances they hold at the central bank for this purpose, or post high quality securities or cash to their account at the central bank. The ability to demand and accept cash and eligible collateral is a necessary function of any viable RTGS system.

1.4 The spread of RTGS systems around the world

Since they emerged in the late 1990s, RTGS systems have become the industry standard for settlement of high value payments. In 1985, only three countries in the world operated an RTGS system. By December 1999, when the Bank for International Settlements (BIS) published the first draft of what became the ten Core Principles for Systemically Important Payment Systems, the number had risen to 25 countries. After the publication of the final version of the Core Principles, the number of countries operating RTGS systems grew exponentially (see Chart 4).

In July 2000, the final version of the BIS Core Principles paper declared, "there has been extensive progress in payment system design in the course of the past ten years, notably in the development and widespread adoption of systems involving real-time gross settlement (RTGS), which can very effectively address the financial risks highlighted by the Core Principles".[3]

Today, the adoption of RTGS systems continues to grow, and has reached 124 systems supporting payments in 160 countries.[4]

3) Committee on Payment and Settlement Systems, Core Principles for Systemically Important Payment Systems, Part 1 – The Core Principles, Report of the Task Force on Payment System Principles and Practices, Paragraph 1.6, page 2, Bank for International Settlements, July 2000.
4) 74 of 125 RTGS systems and 2 of 3 HVP-DNS use SWIFT messaging over the SWIFT network.

"We had always thought that if you wanted to cripple the U.S. economy, you would take out the payment systems. Banks would be forced to fall back on inefficient physical transfer of money. Businesses would resort to barter and IOUs; the level of economic activity across the country would drop like a rock."

Alan Greenspan, The Age of Turbulence, Penguin 2007.

The fact that more countries enjoy the benefits of an RTGS system than there are RTGS systems in existence reflects the fact that several RTGS systems are used by more than one country. Obvious examples include the TARGET2 system operated by the European Central Bank (ECB) in the euro-zone, the shared platform operated by the Banque Centrale des Etats de l'Afrique de l'Ouest (BCEAO) in west Africa, and the equivalent platform operated by the Banque des Etats de l'Afrique Centrale (BEAC) in central Africa.

1.5 The systemic importance of RTGS

The systemic importance of RTGS systems is hard to exaggerate. Any economy which seeks to secure and maintain the confidence of domestic and international investors requires the assurance that payments can always be made, even in the most extreme circumstances. Though RTGS systems are noticeably more important to developed economies and especially major financial centres (see Chart 5), the World Bank has identified the ability of an RTGS system to provide certainty of settlement without credit or liquidity risk as an essential component of the financial infrastructure of any successful economy.

On average, the value of transactions settled by twelve HVPS systems in a group of ten major markets in 2012 was more than 57 times annual national income (as measured by gross domestic product, or GDP).[5]

5) Value of payments processed as a proportion of GDP by the LVTS system in Australia, Canada, Japan, TARGET2 and EBA/EURO1 in the euro-zone, CHATS in the Hong Kong SAR, MEPS+(IFT) in Singapore, RIX in Sweden, SIC in Switzerland, CHAPS in the United Kingdom, and CHIPS and Fedwire in the United States. The data is taken from Bank for International Settlements, Committee on Payment and Settlement Systems, Statistics on payment, clearing and settlement systems in the CPSS countries, Figures for 2012, December 2013.

This ratio varies considerably between countries. In a number of large emerging economies (Brazil, Russia, India, China and South Africa), RTGS payments were in 2012 worth an average of just over 24 times GDP. The ratio was much higher in the STR system in Brazil (38 times) and in the HVP system in China (34 times).

As RTGS systems are adopted by more countries, their systemic importance is increasing. Cross-border transactions mean domestic RTGS systems are also becoming part of a global network of RTGS systems, which in turn links the capital market infrastructures of each country with the capital infrastructures of every country. Domestic PMIs, CSDs and banks are now all part of a complex international eco-system.

In some parts of the world, such as the European Union and west and central Africa, RTGS systems are now formally operating on a regional basis (see Table 1). Some of these regional systems operate from a single shared RTGS platform, while others link a number of separate RTGS platforms. In these regions, it is obvious that the failure of an RTGS system can no longer be confined to one country only. But the same is true of RTGS systems everywhere. They are systemically important, and on a global scale.

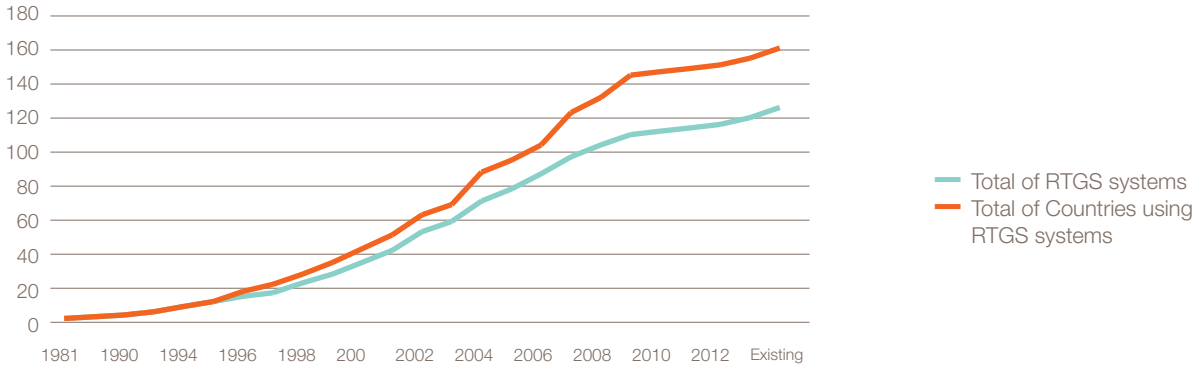

Chart 4: The number of RTGS systems and the number of countries operating an RTGS system

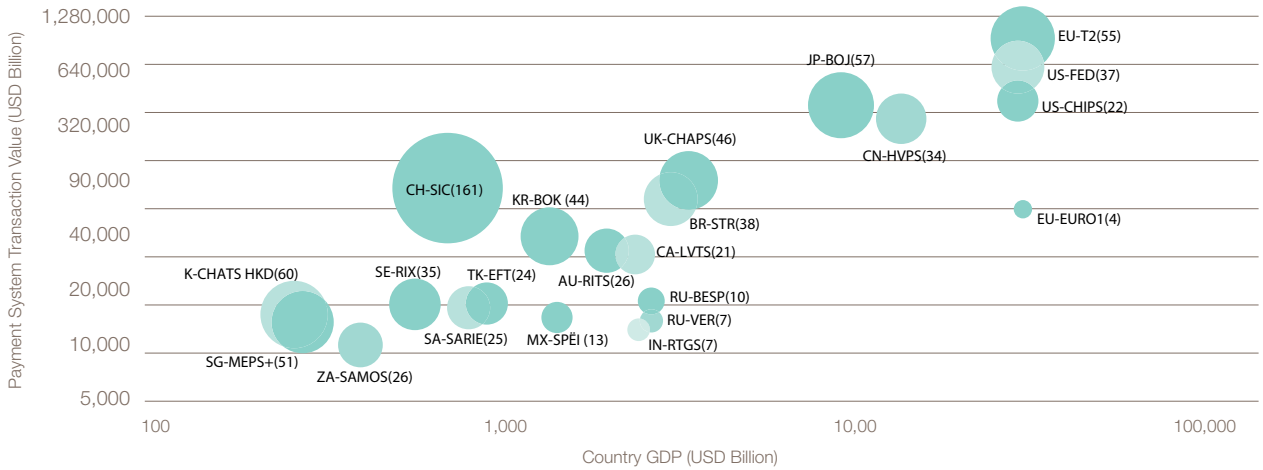

Chart 5: The value of payments processed by RTGS systems as a proportion of GDP. Bubble size is a multiplier of payment transaction value versus GDP. The value after the name of the country and the RTGS system pictures the number of times annual national income are settled by the RTGS over 2012. (Source: Bank for International Settlements, Committee on Payment and Settlement Systems, Statistics on payment, clearing and settlement systems in the CPSS countries, Figures for 2012, December 2013.)

Name of initiative	Model
Association of South East Asian Nations (ASEAN)	In discussion
Banque Centrale des Etats de l'Afrique de l'Ouest (BCEAO)	Single shared platform
Banque des Etats de l'Afrique Centrale (BEAC)	Single shared platform
Common Market for Eastern and Southern Africa (Comesa)	Inter-linked systems
Commonwealth of Independent States (CIS)	In discussion
Consejo Monetario Centroamericano (CMCA)	Inter-linked systems
East African Community (EAC)	Inter-linked systems
Eastern Caribbean Central Bank (ECCB)	Single shared platform
Euro area (Eurozone)	Single shared platform
Gulf Cooperation Council (GCC)	In discussion
Southern African Development Community (SDAC)	Single shared platform
West African Monetary Zone (WAMZ)	Inter-linked systems

Table 1: Regional RTGS systems (Source: SWIFT)

2. Regulatory pressure for resilient RTGS systems

2.1 Threats to RTGS systems

As the systemic importance of RTGS systems has increased, so has the risk that their integrity and security will be compromised. Although none of the threats they face are entirely new – they include the familiar challenges of natural disasters, loss of essential services, data corruption, cyber-attacks, unavailability of staff, component malfunction, terrorism and war (see Chart 6) – their potential to disrupt vital payment services, and their probability of occurrence, have increased significantly in recent years.

Any infrastructure is vulnerable to natural disaster. But when the principal and back-up sites of a PMI are located in the same country or region, the impact of a major natural disaster, such as an earthquake or a flood that disrupts the power grid, or overwhelms the entire geographical area, means all the sites of the RTGS are likely to be out of operation.

Power grids, telecommunications networks and water supplies are all essential services, without which an RTGS system is unable to function. All of these services are vulnerable to disruption by extreme weather conditions, or terrorist attack, but these risks can be mitigated effectively through the provision of alternative sources of supply, or through geographic separation of facilities.

RTGS systems are vulnerable to the failure of technical components. Moreover, this risk is doubled if back-up facilities are based on the same technologies as primary and secondary sites – and almost all are.

The information stored and managed by an RTGS system is subject to data corruption. A malicious attack or an inadvertent error by maintenance staff can cause the corruption or loss of data, or even complete system failure. In addition, as a typical RTGS replicates its data across both its primary and secondary sites, the corruption of data at one site almost invariably leads to the corruption of data at the second site.

As an important part of the financial markets, RTGS systems are an obvious target of cyber-attacks. Deliberate attempts to steal information from RTGS systems, or to take control of their

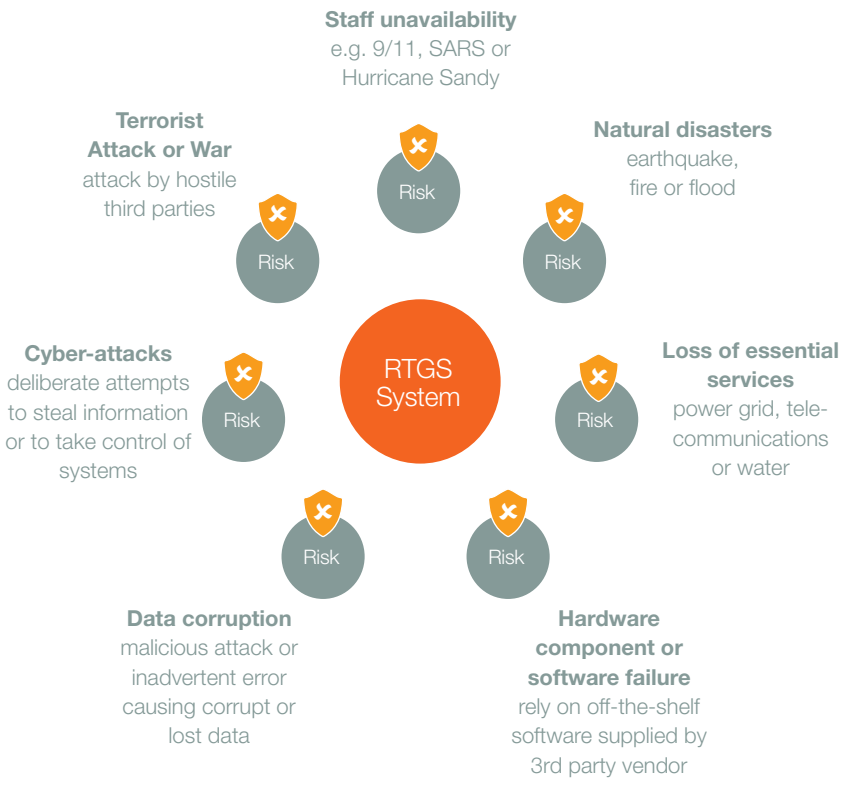

Chart 6: Threats to RTGS systems

systems, are increasing in volume and complexity. They aim to exploit three vulnerabilities: confidentiality, integrity and availability. Attacks on confidentiality aim to steal valuable information. Those on integrity aim to penetrate the systems and alter information or processes without affecting the operation of the system at all. An attack on availability aims to achieve the more straightforward objective of causing the system to stop working. Cyber-attacks are now one of the threats posed by terrorist groups, which aim to intimidate governments or populations by shutting down a crucial national infrastructure.

The strategic profile of RTGS systems makes them ideal targets for attack by hostile third parties, as acts of terrorism or war, either as part of a concerted physical and digital assault, or as a less destructive means of coercing a government or population. In wartime, opposing forces have always fought to destroy or disable infrastructures that are vital to the ability of an enemy country to sustain military, industrial or civilian morale.

Like any organisation, RTGS systems rely on key staff to operate, so the loss or unavailability of employees is a potentially significant threat. Employees may be unable to reach the office because transport networks are down, or choose in difficult circumstances to stay with their families. The Federal Reserve, for example, concluded that the cardinal lesson of 9/11 was the need to disperse its staff and facilities geographically.

Hurricane Sandy, in which the FedWire system operated faultlessly throughout, vindicated that investment.[6]

2.2 The consequences of RTGS failure

If any one of the threats to the security or integrity of an RTGS system is realised, the likelihood is that the ability to send and receive payments will be disrupted. In those circumstances, banks will have to continue to make payments bilaterally, and without having access to liquidity trapped in the disrupted RTGS system. If any bank does not have sufficient liquidity to meet its obligations, it will be forced to go into the market and borrow. In the middle of a crisis, the cost of borrowing is likely to be high and rising.

If a brief disruption to an RTGS system is costly, a prolonged failure would be catastrophic. When banks are unable to makes transfers, the commerce of entire economies slows down, and eventually halts. The money and securities markets, in which governments finance their expenditure, would also slow down and potentially dry up as the reintroduction of credit risk led to a loss of confidence. The equity markets would seize up. Confidence in the economic health of any country affected by the disruption of its RTGS system would evaporate, and its currency would likely collapse.

2.3 RTGS operators are seeking greater resiliency

The damaging consequences of an RTGS system succumbing to any one of a range of plausible threats explains why an appetite for greater resiliency in RTGS systems has grown in tandem with their growing criticality as the guarantor of settlement finality. Resiliency is not the same thing as resilience, which means no more than the ability to recover rapidly from a setback. Resiliency means the ability to continue to operate even if a system has failed completely, by switching activity to a separate system or process or to a collection of separate systems and processes.

Building resiliency of this kind is primarily an objective of the central banks. This is because the overwhelming majority of RTGS platforms settling HVPs are operated by central banks (see Table 2). Understandably, they therefore have the greatest incentive to avoid the potentially catastrophic effects of the failure of an RTGS system, and to contain the consequence of an incident when it occurs. As it happens, several RTGS systems have experienced failure already. However, only the major instances have reached the public domain, and then only because an outage proved impossible to conceal.

Operator	Number of HVP RTGS systems
Central banks	125
Association of commercial banks	3

Table 2: Operators of HVPS systems (Source: SWIFT)

2.4 Regulators demand greater resiliency in all payment systems

As operators of RTGS systems, central banks need no further encouragement to improve the resiliency of the payments infrastructure. However, they and the securities market regulators are also encouraging bank-owned market infrastructures, such as the ACHs, to increase their resiliency. The 24 CPSS-IOSCO principles published in April 2012, drawn up by the international co-ordinating bodies of the central banks and the securities markets regulators, are a set of best practice recommendations that span both payments and securities market infrastructures. CPSS-IOSCO expects them to be implemented at the domestic level.

The fact that many of the CPSS-IOSCO Principles address some aspect of settlement is an indication of the importance international regulatory bodies now attach to the achievement of settlement finality in central bank money as promptly as possible. Indeed, the Principles are based on the assumption that final and irrevocable settlement of transactions in real time and in central bank money will become ever more widely available, including through ACHs, CSDs as well as RTGS systems.

The corollary of that ambition is the emphasis in the Principles on greater resiliency. Principle 17 (see Table 3) sets the highest standards for FMI's in terms of security, operational reliability, scalability and business continuity, including rapid recovery of service in the event of disruption. Indeed, Principle 17 recommends that any disruption last no more than two hours, and that normal service should be resumed before the end of the business day on which the disruption occurs, even in the case of a catastrophic situation that impacts the primary operational sites of the FMI.

The pressure is therefore on FMI's, which includes PMIs operating RTGS styem, to improve their contingency arrangements in order to ensure business continuity. To satisfy the detailed demands of Principle 17, an operational contingency plan must deliver continuity of service through the most improbable of catastrophic events (called extreme circumstances in the Principle 17) as well as marginal or critical failures (called disruptive events in the Principle 17).

2.5 The cost of building a resilient system

The necessary degree of resilience will not be purchased cheaply. RTGS systems are expensive to develop in the first place, because they have to be impregnable, though the price is rarely unaffordable. It is the cost of making the RTGS system sufficiently resilient enough to withstand technical failure, natural disasters, loss of essential services such as electricity, terrorist attacks, sabotage, data corruption and cyber-attacks that adds substantial costs. At present, these are inhibiting the construction by RTGS systems of operational contingency and disaster recovery plans sufficiently credible to satisfy the standard set by Principle 17 of the CPSS-IOSCO Principles.

6) Richard P. Dzina, wholesale product office, Federal Reserve, quoted in "Failure not an option," in MI Forum magazine, Sibos Dubai, 2013, page 18.

A comparison of the 2000 CPSS and 2012 CPSS-IOSCO principles for payment systems

10 core principles for payments systems only, finalised by CPSS in July 2000	Relevant principles from 24 applied to all financial market infrastructures (FMIs) by CPSS-IOSCO in April 2012
The system should have a well-founded legal basis under all relevant jurisdictions.	**Principle 1:** Legal basis: An FMI should have a well-founded, clear, transparent, and enforceable legal basis for each material aspect of its activities in all relevant jurisdictions.
The system's rules and procedures should enable participants to have a clear understanding of the system's impact on each of the financial risks they incur through participation in it.	**Principle 3:** Framework for the comprehensive management of risks: An FMI should have a sound risk-management framework for comprehensively managing legal, credit, liquidity, operational, and other risks. **Principle 23:** Disclosure of rules, key procedures, and market data: An FMI should have clear and comprehensive rules and procedures and should provide sufficient information to enable participants to have an accurate understanding of the risks, fees, and other material costs they incur by participating in the FMI. All relevant rules and key procedures should be publicly disclosed.
The system should have clearly defined procedures for the management of credit risks and liquidity risks, which specify the respective responsibilities of the system operator and the participants and which provide appropriate incentives to manage and contain those risks.	**Principle 4:** Credit risk: An FMI should effectively measure, monitor, and manage its credit exposures to participants and those arising from its payment, clearing, and settlement processes. **Principle 5:** An FMI that requires collateral to manage its or its participants' credit exposure should accept collateral with low credit, liquidity, and market risks. An FMI should also set and enforce appropriately conservative haircuts and concentration limits. **Principle 7:** Liquidity risk: An FMI should effectively measure, monitor, and manage its liquidity risk. An FMI should maintain sufficient liquid resources in all relevant currencies to effect same-day and, where appropriate, intraday and multiday settlement of payment obligations with a high degree of confidence under a wide range of potential stress scenarios
The system should provide prompt final settlement on the day of value, preferably during the day and at a minimum at the end of the day.	**Principle 8:** Settlement finality: An FMI should provide clear and certain final settlement, at a minimum by the end of the value date. Where necessary or preferable, an FMI should provide final settlement intraday or in real time. **Principle 9:** Money settlements: An FMI should conduct its money settlements in central bank money where practical and available. If central bank money is not used, an FMI should minimise and strictly control the credit and liquidity risk arising from the use of commercial bank money. **Principle 10:** Physical deliveries: An FMI should clearly state its obligations with respect to the delivery of physical instruments or commodities and should identify, monitor, and manage the risks associated with such physical deliveries. **Principle 12:** Exchange-of-value settlement systems: If an FMI settles transactions that involve the settlement of two linked obligations (for example, securities or foreign exchange transactions), it should eliminate principal risk by conditioning the final settlement of one obligation upon the final settlement of the other. **Principle 20:** FMI links: An FMI that establishes a link with one or more FMIs should identify, monitor, and manage link-related risks.

A system in which multilateral netting takes place should, at a minimum, be capable of ensuring the timely completion of daily settlements in the event of an inability to settle by the participant with the largest single settlement obligation.	**Principle 13:** Participant-default rules and procedures: An FMI should have effective and clearly defined rules and procedures to manage a participant default. These rules and procedures should be designed to ensure that the FMI can take timely action to contain losses and liquidity pressures and continue to meet its obligations.
Assets used for settlement should preferably be a claim on the central bank; where other assets are used, they should carry little or no credit risk and little or no liquidity risk.	**Principle 15:** General business risk:An FMI should identify, monitor, and manage its general business risk and hold sufficient liquid net assets funded by equity to cover potential general business losses so that it can continue operations and services as a going concern if those losses materialise. Further, liquid net assets should at all times be sufficient to ensure a recovery or orderly wind-down of critical operations and services. **Principle 16:** Custody and investment risks: An FMI should safeguard its own and its participants' assets and minimise the risk of loss on and delay in access to these assets. An FMI's investments should be in instruments with minimal credit, market, and liquidity risks.
The system should ensure a high degree of security and operational reliability and should have contingency arrangements for timely completion of daily processing.	**Principle 17:** Operational risk: An FMI should identify the plausible sources of operational risk, both internal and external, and mitigate their impact through the use of appropriate systems, policies, procedures, and controls. Systems should be designed to ensure a high degree of security and operational reliability and should have adequate, scalable capacity. Business continuity management should aim for timely recovery of operations and fulfilment of the FMI's obligations, including in the event of a wide-scale or major disruption. An FMI should aim to be able to resume operations within two hours following disruptive events; however, backup systems ideally should commence processing immediately. The plan should be designed to enable the FMI to complete settlement by the end of the day even in case of extreme circumstances.
The system should provide a means of making payments which is practical for its users and efficient for the economy.	**Principle 22:** Communication procedures and standards: An FMI should use, or at a minimum accommodate, relevant internationally accepted communication procedures and standards in order to facilitate efficient payment, clearing, settlement, and recording.
The system should have objective and publicly disclosed criteria for participation, which permit fair and open access.	**Principle 18:** Access and participation requirements: An FMI should have objective, risk-based, and publicly disclosed criteria for participation, which permit fair and open access. **Principle 19:** Tiered participation arrangements: An FMI should identify, monitor, and manage the material risks to the FMI arising from tiered participation arrangements.
The system's governance arrangements should be effective, accountable and transparent.	**Principle 2:** Governance: An FMI should have governance arrangements that are clear and transparent, promote the safety and efficiency of the FMI, and support the stability of the broader financial system, other relevant public interest considerations, and the objectives of relevant stakeholders.

Table 3: Source: Committee on Payment and Settlement Systems, Core Principles for Systemically Important Payment Systems, Bank for International Settlements, July 2000, and Committee on Payment and Settlement Systems (CPSS) and Technical Committee of the International Organisation of Securities Commissions (IOSCO), Principles for Financial Market Infrastructures, Bank for International Settlements, April 2012.

3. Best practices in RTGS resiliency planning

3.1 A second operational site

Continuous availability of real-time payment services requires a high degree of resiliency. The first step towards achieving it is to build a complete facsimile of the primary RTGS platform at a different location. This covers the risk that the primary RTGS site fails. It is a sufficiently unavoidable measure to ensure that RTGS systems in all major financial markets are now supported by a second site.

Indeed, countries installing an RTGS platform for the first time now include the construction of a back-up site as a matter of routine. However, even if it has become best practice, a second site is not yet a standard feature in small or medium-sized countries with an established RTGS system.

A second site back-up has to take over from the primary site seamlessly. As Principle 17 of the CPSS-IOSCO Principles prescribes, a second site must be capable of activation no more than two hours after a relatively commonplace event causes the failure of the primary platform. The World Bank survey of 2010 indicates an anticipated delay in RTGS business continuity planning scenarios for full system recovery of anywhere between ten and 240 minutes, with an average of 100 minutes.[7]

In practice, cut-over is achieved much more rapidly than these estimates imply, since most second sites are run on so-called "hot standby." This condition means the primary and secondary platforms operate in parallel, with the back-up site replicating the work of the first site in a synchronised way. This enables the second site to take over from the first site more or less immediately.

For "hot standby" to work faultlessly, rapid replication at the secondary site of all information being processed by the primary site is essential. If information is not completely up-to-date at the time of the failure of the primary site, disputes will arise between banks over whether a transaction has settled or not.

The CPSS-IOSCO principles assume that a second site will cover all the probable and frequent threats to the integrity of an RTGS platform that are likely to arise. The implicit corollary is that a second site cannot cover the infrequent or improbable threats, and cannot cover the total loss of both sites.

3.2 Reversion to bilateral arrangements

If both the primary and secondary sites of an RTGS system are lost, the simplest solution is to revert to bilateral settlement arrangements, by which all payments are exchanged directly between the banks on a net basis at the end of the business day and not transaction by transaction on a gross basis via the RTGS system.

However, bilateral settlement presents a number of difficulties. One RTGS operator predicts that bilateral payments would have to be undertaken by a combination of manual and automated procedures. It estimates that this would require an additional 15-20 permanent staff at a remote centre, which is the minimum number needed to cope with an average of 100,000 bilateral transactions a day.

Secondly, as full service recovery needs accurate records of payments made before as well as after the point of failure, a prolonged dependence on manual processing of bilateral settlements will be time consuming and complicated to control. In particular, it puts the onus on the participating banks to identify the highest priority payments, and account for their despatch and receipt.

Thirdly, even the manual processing of bilateral settlements depends on a minimum set of RTGS functionalities remaining available to settle the multilateral net positions at the end of each business day. If both primary and secondary sites are lost, this is unlikely to be the case.

Last but not least, the replacement of settlement through an RTGS system by bilateral arrangements reintroduces credit risk into the settlement process, not least because multilateral net positions are settled at the PMI at the end of the day through complex procedures that rely on the RTGS operator being available at the end of the day to settle, which would not necessarily be the case.

Only immediate settlement finality can eliminate credit risk, and bilateral settlement processing makes that impossible.

3.3 A third operational site

To overcome these disadvantages, some RTGS platform operators have built a third back-up site in a geographically remote location. For instance, a handful of major economies (the European Union, the United States, Russia, China, India and Brazil), all of which boast significant populations, geographical diversity and size, as well as sufficient transactional activity, have built third sites. For most, however, building a third site, even in a different part of the same country, runs the risk of adding little to the threat coverage already provided by the second site, whilst increasing costs to potentially prohibitive levels.

To achieve total differentiation is extremely expensive. Estimates of the cost of operating a third site suggest a minimum of 10 to 15 percent of the costs of the primary and secondary sites, even without the benefits of geographical and staff differentiation. It is clear that a fully differentiated third site demands substantial investment.

3.4 The need for speed in resuming an RTGS service

In most circumstances, the CPSS-IOSCO Principles expect a back-up RTGS system to resume settlement within two hours. In the extreme circumstance in which both back-up sites have failed, the CPSS-IOSCO Principles allow RTGS payments to resume settlement operations by the end of the day of the disruption. After all, failure to do so would reintroduce credit risk, with the attendant possibility that confidence in the financial system unravels completely.

Importantly, the next-day timetable is more demanding than it sounds. If both sites are lost at, say, 9.00 am, the RTGS system has until the end of the trading day to recover. If it is later in the day - say, 3.00 pm – there is much less time to reactivate the primary site or activate the back-up site before transactions resume the following morning.

7) World Bank Global Payment Systems Survey 2010, Appendix, Table II.15.

Cross-currency settlement, which requires the exchange of principal in two currencies, imposes another time constraint. In major currency pairs, banks send settlement instructions to CLS Bank, which simultaneously settles netted payment instructions across accounts maintained by the banks at CLS. This can only occur during a five hour window when the RTGS systems in the 17 currencies settled by CLS Bank are open.

3.5 Capturing balances and transactions at the point of failure

To take over as rapidly as the CPSS-IOSCO timetable implies, a back-up site needs accurate information about balances to restart. In other words, knowing which payments have settled and which are pending. Once accurate starting balances are established, the system must be able to effect transfers between banks on a gross, transaction by transaction basis in real time, including the settlement of net payments between banks agreed via ACHs.

If the balance information is available online, the time-lapse before re-starting should be minimal, but ensuring the security and integrity of information online can be extremely expensive. Storing the necessary information offline is cheaper, but it means the back-up solution cannot start with the latest information that was being processed at the point of failure, because the data is not transferred to the back-up site in real time. The resultant gaps in data have to be covered by dialogue with and between the banks.

Either way, rapid restoration of service is possible only on the basis of accurate information about the balances banks hold at the central bank, and which transactions have settled and which are pending at the point of failure of the RTGS system. This is the only way to avoid disputes between banks over what was settled versus what was unsettled, which would delay recovery from disruption.

To work successfully, a back-up solution needs more than the balances outstanding. It also needs to host configuration data – such as the account structure – of the failed RTGS system, in order to resume payments immediately. It should also offer system operators online reporting tools that allow them to view, understand and manage the status of their transactions and their liquidity positions accounts, and set the credit limits that prevent participating banks becoming overdrawn at the central banks.

3.6 Coverage of existing business during the recovery period

Depending on the scale of the failure, market participants may have to continue to work with the back-up service for a number of days, and possibly a number of weeks or months. During this time, the back-up platform has to be readily accessible to all the banks that were using the failed RTGS system, and also be capable of handling the same volume and value of transactions.

Limiting the number of participants and/or the number of transactions serviced by a back-up service is one way of limiting the cost of setting of an operational contingency plan, but it reintroduces credit risk in transactions processed outside the system. Best practice argues unequivocally for full coverage of all existing participants and transactions.

This is an unavoidable constraint on any operational recovery plan. A back-up solution which lacked the capacity to process the same amount of business as the status quo ante the failure of the primary RTGS system would force the central bank to oblige at least some counterparties in the market to settle bilaterally. This would reintroduce a degree of credit risk.

3.7 Necessity of minimal impact on the users

An operational back-up service capable of capturing the data to resume payments quickly, and of processing high volumes without reintroducing credit risk, sets demanding criteria. Yet there is a further demand the service has to meet: minimal impact on the day-to-day operational activities of the banks that make use of the service.

Any back-up system which denies its users vital information, or forces them to make a major investment in new technology or additional staff, or whose shortcomings prompts them to question its reliability, will constrain the effectiveness of the third line of defence. This argues for selecting, tried and tested infrastructures with which users are already familiar.

3.8 The value of diversity

Geographical remoteness can add greatly to the resilience of a back-up facility, especially if the facility is located in another country where it is physically removed from the centre of a natural disaster or terrorist attack, and linked to entirely separate sources of energy and communications.

However, most back-up sites have tended to be built in fairly close proximity to the primary site. This is largely because the second site has to capture and synchronise information in near real time, and its ability to do so is impaired by distance. Current data transmission speeds restrict the effective distance to between 50 and 100 kilometres.

To achieve the requisite degree of geographic remoteness, a third site must sacrifice some information about payments completed and pending at the point of failure, either because of the natural latency of online data transmissions or because the data was stored offline, in batch mode.

Technical diversity is another important form of defence, providing a further layer of protection against discrete risks - viruses, component failure, data corruption, human error, malicious insiders and cyber-attacks. Outsourcing the second (or third) back-up to a separate and wholly independent operator with distinct dependencies reducess the associated risks.

Last but not least, reducing reliance on locally based staff further increases resiliency. Loss of essential services, a natural disaster or a terrorist attack affecting a wide geographical area is bound to prevent at least some staff getting to work. A back-up facility in a separate location that employs different people, or which allows the same people to operate it remotely when they are unable to reach their normal place of work or first standby facility, overcomes these risks.

3.9 Independent data storage

Adhering to all of these best practices may still be insufficient to make an RTGS system truly resilient. A back-up system can be geographically and technically differentiated, capable of re-starting quickly with the latest transaction status reports and liquidity balances, impose minimal requirements on existing users, and be operated by a separate cadre of staff, yet still be vulnerable to a cyber-attack through the transmission of corrupted information from the primary or secondary site. That means the third site will re-start payments on the basis of incorrect information.

The only solution to this risk is the independence of the back-up site from the primary or secondary site. Independence entails the safe storage of the data by an independent and trusted third party organisation on a continuous basis in real time, so that the latest balances are always retrievable, along with a full explanation of how those balances were derived, together with a full archive of transactions between the counterparties.

In the event that recent information is corrupted, the archive facilitates the rewinding of transactions beyond the point at which the data was corrupted to ascertain the correct balances. A contingency site can then be confident of re-starting the payments process with a so-called "golden copy" of uncorrupted and accurate information about the status of balances and transactions at the point of failure. A back-up site can use this "golden copy" to re-commence settlements at the point of the original failure.

3.10 Pooling resources

All of these best practices (see Table 4) add to the complexity of managing the risk of failure of an RTGS system. If the third line of defence is based in another country, operates on separate technology to different processes and procedures, and relies on independent storage of balance and transaction data, it is inevitably harder and more expensive to manage the various elements that make up the system as a whole.

It follows that, to the extent that the challenges of maintenance and management of a back-up facility can be shared, the costs of that additional complexity will be reduced.

> Back-up platform to operate throughout a crisis without reintroducing credit risk
> Resume the RTGS service with a speed that avoid re-introduction of credit risk
> Capture the latest balances and transaction status continuously in real-time
> Avoid any restriction on the number of participants and volumes of transactions
> Minimise the impact on users of the switch to the back-up system
> Ensure the back-up facility is geographically remote from the other sites
> Technical diversification will limit the risk of cross-contamination
> Reduce reliance on essential staff to operate systems
> Independent, trusted storage of balances and messages exchanged between counterparties

^ Table 4: Best practices in RTGS risk mitigation

However, a service in which a number of RTGS systems pool their resources can do more than save time and money. It can also facilitate the adoption and development of best practices.

4. The solution

4.1 The lack of contingency solutions that meet best practices

Best practices in RTGS back-up systems include a rapid cut-over to the service; a geographically remote facility; reduced reliance on local staff; technical diversity; independent data storage; sufficient capacity to support existing volumes of business; minimal impact on users; the availability of the service throughout the period of disruption; and, most importantly of all, the ability to capture a clear view of the intra-day balances at the point of failure, or recreate it rapidly once the primary and secondary sites have failed. Understandably, these requirements are hard to meet at reasonable cost.

4.2 How SWIFT can help

In 2009, SWIFT started to review the feasibility of achieving exactly that: a back-up RTGS system that could be made available to RTGS operators at reasonable cost, because it was a shared service. As it happens, the Bank of England had embarked simultaneously on an investigation of how it could further increase the resilience of its own RTGS system.

Like every central bank, the sensitivity of the Bank of England to the systemic risk created by the reliance of the British economy on RTGS was heightened by their experience of the acute phase of the financial crisis between 2008 and 2009. Although no RTGS system failed during the crisis, the episode reminded all central banks that the failure of even one RTGS system would accelerate and aggravate a crisis, because transfers of cash and other assets would be impeded.

The Bank of England was concerned enough to explore how payments could continue to settle if both its primary and secondary sites were disabled. Its first assumption was that banks could revert to bilateral settlement.[8] Unfortunately, this assumption encountered two obstacles.

The first was that the banks which use the Bank of England RTGS system were unable to handle their existing volumes of bilateral payments on a net basis once the central bank was reduced to reliance on manual processing. If the back-up system could not provide sufficient capacity to solve that problem, credit risk would be reintroduced.

A second obstacle was that no back-up system could start processing, even on a manual basis, without agreement on the starting balances. No solution could be found to ensure that the Bank of England had access to payments balances at the point of failure.[9]

Even if solutions to these two obstacles could be found, the Bank of England quickly concluded that the costs were hard to justify in relation to the benefits. So it welcomed the opportunity to work with SWIFT and six other central banks on the feasibility of SWIFT hosting an RTGS payments system.

The involvement of a sizeable group of central banks ensured that the service was sufficiently generic for it to be applied across a wide variety of RTGS models. By early 2011, the Bank of England was ready to serve as a pilot site. Following formal agreement between SWIFT and the Bank of England, implementation began in earnest.

Testing of the technology and processes began on 29 July 2013, and the new system, called “Market Infrastructure Resiliency Service” (MIRS), went live with the Bank of England on 24 February 2014.

4.3 What is the Market Infrastructure Resiliency Service (MIRS)?

MIRS is a generic RTGS system available to central banks that use the SWIFT network. It is designed to replace the functionality essential to achieve final, irrevocable settlement of payments on a transaction by transaction basis in real time on behalf of any RTGS system (see Table 5).

However, MIRS can also be customised to meet the specific RTGS needs of any contingency plan, including features that are not part of the generic functionality of the service.

Through MIRS, SWIFT believes that it can play a pivotal role in supporting central banks in their efforts to protect themselves against large-scale failures of their RTGS systems. MIRS satisfies the best practices in RTGS back-up systems too, including speed of transition from the failed RTGS systems; geographical remoteness; a separate group of operational staff; completely differentiated technology; independent and trustworthy data storage; ample capacity to handle current volumes of transactions; continuous availability throughout the period of disruption; and especially the ability to capture a clear view of the intra-day balances at the point of failure.

MIRS is a generic 3rd site back-up solution
MIRS enables best practice resiliency in accordance with the new principles for FMI
MIRS is technologically and geographically diverse from the existing RTGS system
MIRS is a shared service, making the solution cost efficient
MIRS reconstructs balances at point of failure based on data in an safe store
MIRS is capable of processing and settling payment transactions, on a high-capacity basis
MIRS is activated and deactivated by the RTGS operator
MIRS can run as long as it is needed

Table 5: What MIRS does

8) See 3.2 Reversion to bilateral arrangements

9) See 3.5 Capturing balances and transactions at the point of failure

SWIFT also believes that MIRS can provide the back-up service at a relatively low cost by comparison with building a third RTGS back-up site in a separate location using different staff and technology, and with minimum impact on all market participants, including ACHs, direct participants, indirect participants and CSDs (see Chart 7).

4.4 The principal benefits of MIRS

The most important of the benefits of MIRS is its ability to capture payments transaction balances on a regular basis, and so deliver a clear view of the position at the point of failure of the primary and secondary RTGS sites. It achieves this by making use of the settlement confirmation messages stored in the SWIFT databases in the normal course of business.

By capturing data continuously, MIRS ensures its service availability within 2½ hours of the request for its activation, eliminating credit risk. It is also built with sufficient scale to absorb the full transaction volume processed by any existing RTGS system, and is robust enough to continue to operate throughout the period required to restore the primary and secondary services, whether the recovery period is measured in days or weeks or months.

By shifting processing to a SWIFT facility, MIRS has the further benefit of increasing geographical diversity. Since it also depends on SWIFT technical platforms and messaging, all of which are in-house developments rather than purchases from RTGS application vendors, MIRS, coupled with SWIFT's financial messaging services, also guarantees technological diversity, increasing the resilience of the RTGS to cyber-attack or any other form of data or process corruption.

Because MIRS is a service hosted by SWIFT, switching to MIRS has a minimal impact on users, since most banks and RTGS system operators are already members of the SWIFT network, and routinely make use of its infrastructure and message types, especially in cross-border payments.

As RTGS operators, the central banks that make use of MIRS in an emergency do not have to master a range of new technical interfaces and operational techniques. The system is designed to be operated by a minimal number of staff, using tools that

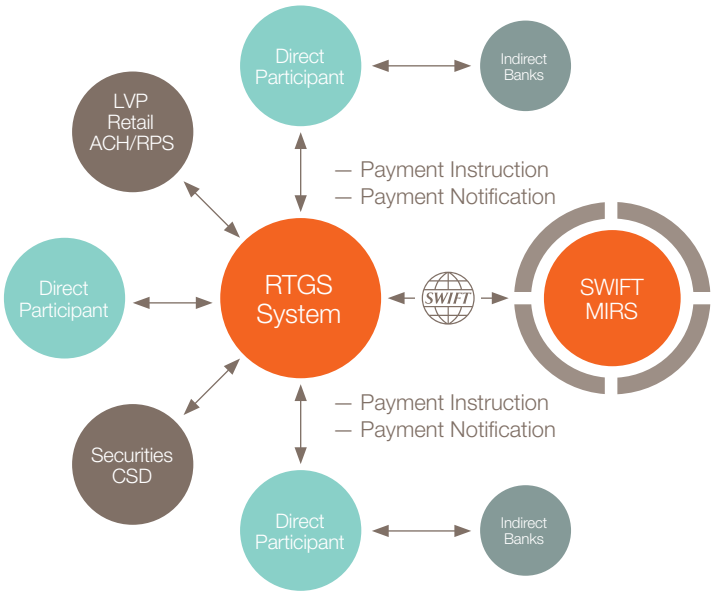

Chart 7: How MIRS backs up RTGS systems

free them to concentrate on maintaining operations from a business perspective only, and not from a technical point of view.

Lastly, MIRS is extremely secure. It makes use of the existing infrastructure of SWIFT, which operates to the highest global standards in terms of security and resiliency. Although MIRS is a service that can be shared by multiple RTGS systems, there is no risk of information leakage between them because each user accesses a fully segregated service.

4.5 RTGS operators are always in control of MIRS

Importantly, SWIFT will at no point actually be operating MIRS. Once activated, the operators of the failed RTGS system remain in full control of all payments transactions, with SWIFT acting as host of the system only. There is no possibility that SWIFT will initiate transfers between accounts, alter credit limits or respond to business alerts. In addition, there is no possibility that SWIFT will activate or deactivate the system, as for security reasons this action must be managed by the RTGS operator.

RTGS operators control their interactions with MIRS either through secure web access from a "command centre" established by the central bank as operator of the RTGS system.

Naturally, the web service offers richer functionality to the RTGS operators, which have to open accounts, set credit limits, manage transaction queues and relieve gridlocked payments on behalf of all participants. RTGS system participants, on the other hand, need only to initiate, view and manage their own profile, account and payments (see Table 6).

The functionality offered by MIRS to users was agreed in consultation with seven separate central banks, and is generic across all of the different RTGS systems they operate. This affords a high degree of confidence that the functionality of MIRS can be adapted successfully for use by most RTGS systems. This is important, given that MIRS will be providing support for a wide variety of RTGS systems, and almost always in a period of crisis, when prompt switching of services will be essential.

It is of course conceivable that a crisis is so extreme that an RTGS operator is unable to operate MIRS. In such an extreme situation, the RTGS operator is free to make arrangements with a trusted party – such as another central bank that operates an RTGS system (a "buddy RTGS Operator") – to activate and operate MIRS on their behalf.

Functions	RTGS Operator (e.g. Central Bank)	RTGS participants (e.g. Banks, ACH,CSD,…)
Configuration data	Create/Update	View
System state information	View	n/a
Windows	Create/Update	n/a
Business calendar	Create/Update	View
Exchange rates	Create/Update	View
Checkpoint status	View	n/a
Activation	Initiate	n/a
Deactivation	Initiate	n/a
Intra-day account information	View	View
Payment Information/Queue management (Transation log)	Update	Own payments
Alerts	Manage	View
Liquidity parameters (limits)	Manage	View
Account balances SoD and EoD	Create/View	View
On-line transfert (Manual or file upload)	Create/Upload	Create/Upload
Suspend/resume settlement	Update	n/a
Gridlock	Initiate	n/a
Community messages	Create/View	View/Create to MI
Audit information	View	View
MIRS archives	Download	n/a

Table 6: Functionality available to RTGS operators and participants

It is entirely possible that two RTGS operators, who both use MIRS, will come to a mutual arrangement to operate MIRS on behalf of each other in the event that either is unable to do so independently.

4.6 MIRS can be activated quickly

MIRS aims to move from its dormant mode to active mode in less than 2½ hours of the request for its activation (see “MIRS operational modes”). It can achieve this provided the latest snapshot of the balance position dates back no further than 55 minutes prior to the failure of the primary RTGS system. Once MIRS is activated, all types of participant, and transactions of any volume and value, can be supported in their entirety.

MIRS operational modes

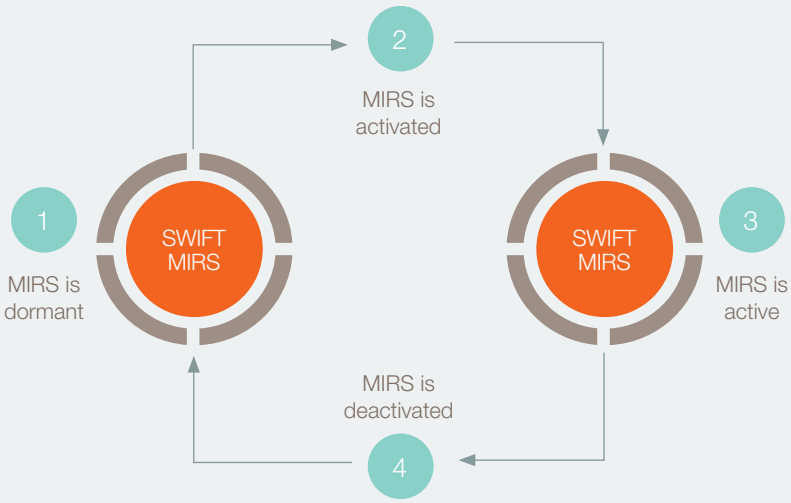

Chart 8: MIRS Operational Modes

MIRS operates in either dormant mode or active mode, and switches between the two states as it is required to support a disabled RTGS system.

By default, MIRS operates in dormant mode, awaiting activation. Once an RTGS system operator requests activation and confirms it, MIRS switches to active mode, and continues in that state until the operator requests deactivation. The business operation of the payment system remains with the RTGS. Chart 8 illustrates how it moves from dormant to active mode and back again.

When MIRS is **dormant**, payment instructions and related notification messages are exchanged between the RTGS and participants using SWIFT.

SWIFT stores all messages in its FIN database, from which they can be retrieved when MIRS is activated. This means that SWIFT systems, which are already being used by the majority of global RTGS platforms, are automatically capturing payment instructions and settlement notifications between participants continuously in real time, and storing them in their databases. Since copies of the payment messages are stored by SWIFT before they are processed by the RTGS platform, where they might be corrupted or lost, their validity and accuracy is also assured. These stored messages provide the starting point for MIRS in activation mode.

In addition to the automatic storage of the individual payment messages by SWIFT, the RTGS is obliged to send MIRS a time-stamped balance of each of its accounts, known as a "balance checkpoint", every 55 minutes, or less. These snapshots of the account balances will be sent by the RTGS using the SWIFT balance checkpoint message type. MIRS will validate the content of the balance checkpoint message to ensure configuration coherence. The checkpoint balances, together with the settlement confirmation messages, will be used by MIRS when it is called upon to take over from a failed RTGS platform.

As MIRS stores balance checkpoint messages for the last four days, MIRS will be able to re-start payment settlement services at any point in time within the last four days, when requested.

RTGS platforms that do not use SWIFT messaging would be obliged, after selecting MIRS as a back-up service, to send to SWIFT copies of settlement confirmation messages and balance checkpoint messages in the SWIFT message format, over the SWIFT network. This necessitates additional but manageable effort.

Finally, in dormant mode, MIRS provides the RTGS operators with secure web access for the maintenance of participant and account static data, alerts, monitoring functions, calendar.

Following the failure of its operational sites, the RTGS system operator requests **activation** of the MIRS service.

Upon receipt of this request, which requires "four eyes" validation, double initial secret entry by operators and "four eyes" re-confirmation, MIRS assumes the Bank Identifier Code (BIC) of the central bank that operates MIRS. The service will retrieve all the relevant payment and settlement confirmation messages from the SWIFT database that occurred after the balance checkpoint to the time of the failure.

By starting from the account balance checkpoint selected by the RTGS operator, MIRS verifies which payments have already settled versus FINCopy payments that are still pending. This process enables MIRS to provide the RTGS operator and its participants with the correct account balances prior to the point of failure, as well as the FINCopy outstanding payments to be queued for settlement. MIRS will also ensure that legitimate settlement confirmations that have not been notified to the relevant participants are identified.

Typical examples are failures that occur after the RTGS platform settled payments but before or during the sending of the notifications to the relevant participant or failures when not all notifications resulting from gridlock resolution or from failed liquidity optimisation cycles.

In any event, once an accurate starting position has been determined, MIRS provides a status report to the RTGS operator, which then has the option to proceed with the activation or cancel it.

Confirmation by the RTGS operator will transition MIRS into **active** mode.

Once activation is confirmed by the RTGS operator, MIRS takes the role of the failed RTGS system. Participants will continue to send payment instructions over SWIFT, as normal, and these payments will be routed automatically to MIRS, with MIRS then assuming the role of the RTGS.

From its starting position, missing notifications identified during the activation will be sent guaranteeing the community a correct reconciliation between the latest accounts balances used by MIRS and the account information hold by the participants. MIRS will have determined the payments that need to be queued for settlement. These transactions and new payment instructions are then processed in real time, transaction by transaction, and as normal. MIRS will act as the RTGS, settling or rejecting the transactions, using the pre-agreed credit rules, defined by the RTGS.

MIRS will receive and validate each message against a number of criteria (such as whether the account and currency code and business day are valid, and not a duplicate, and that the sender is authorised to act on behalf of the account).

Once this validation process is complete, the payments will be submitted for settlement. After checking that there are sufficient funds are available in the account to meet them and that no other payments with a higher priority need to be settled before them, MIRS will settle each payment by simultaneously debiting and crediting the accounts of the direct or indirect participants. If required, it sends notifications to both.

As part of the settlement process, MIRS will also provide various management tools for the RTGS. These include tools for monitoring and control purposes (such as the state of the system, account monitoring, alerts, and queue monitoring) and for the management of liquidity and reporting.

Once the primary RTGS system is repaired and in a position to resume operation, the operator requests **deactivation** of MIRS. As the deactivation process is a planned process, it will take place outside of business hours. This request is subject to "four eyes" validation.

Upon deactivation, MIRS will provide its balance checkpoint message back to the repaired RTGS, plus a copy of its archive which contains both audit log and transaction log. The balance checkpoint will be used by the RTGS as its new starting position as it resumes operation. MIRS then deactivates itself, and the RTGS system resumes control of its operations. The participants continue to send payments as usual.

Legal notices

About SWIFT
SWIFT is a member-owned cooperative that provides the communications platform, products and services to connect more than 10,000 financial institutions and corporations in 212 countries and territories. SWIFT enables its users to exchange automated, standardised financial information securely and reliably, thereby lowering costs, reducing operational risk and eliminating operational inefficiencies. SWIFT also brings the financial community together to work collaboratively to shape market practice, define standards and debate issues of mutual interest.

Disclaimer
SWIFT supplies this publication for information purposes only. The information in this publication may change from time to time. You must always refer to the latest available version.

For more information about SWIFT, visit www.swift.com.

57023 - MAY 2014