A Broadcom Company
technologies

Solution Brief

CA Dynamic Capacity Intelligence

Intelligent Mainframe Capacity Management

Key Benefits

- **Optimize system resources for prioritized workloads.** CA Dynamic Capacity Intelligence, for z Systems, provides proactive, predictable capacity management so that you can focus on what is most important.
- **Automate intelligently.** Avoid unplanned spikes in cost and empower IT to better manage service level agreements (SLAs) through automated, dynamic capacity optimization.
- **Improve management of SLAs.** Enable continuous optimization of mission critical application delivery and operations.
- **Automate continuous optimization for priority workloads.** Shift available mainframe capacity across LPAR boundaries.
- **Contain and control software costs.** Flexibility to help provide that critical workloads complete per SLAs with minimal impact to MLC.

Overview

CA Dynamic Capacity Intelligence from CA Technologies is an easy-to-use, fully automated, intelligent capacity management solution. The solution is designed to streamline your ability to monitor and effectively manage SLAs and to improve operational efficiency. CA Dynamic Capacity Intelligence helps IT to achieve rapid value by enabling priority mainframe workloads, meeting SLAs, and removing the guesswork from MLC pricing.

Business Challenges

For many businesses, any breach of service level agreements (SLAs) for priority workloads is unacceptable. The dilemma is that IT must provide the best possible service to the business, while the business must prioritize cost containment strategies.

Opportunity

IT and business stakeholders desire a predictable and proactive way to ensure that critical workloads complete in the time limits that are specified by SLAs. At the same time, business stakeholders want a minimal impact to cost. Selecting the best possible combination of pricing options, monitoring variables, and resource management strategies can lead to significant savings.

Solutions Overview

CA Dynamic Capacity Intelligence for IBM z/OS systems provides proactive, predictable capacity management for optimizing system resources and prioritizing workloads. CA Dynamic Capacity Intelligence helps you to better manage and use mainframe capacity, driving down monthly license charges (MLC). Intelligent automation avoids unplanned spikes in cost and empowers IT to better manage service level agreements (SLAs) through automated, dynamic capacity optimization.

CA Dynamic Capacity Intelligence helps you to achieve all of the following goals:

Take the Guesswork out of MLC Pricing
IT must provide the best possible service to the business. At the same time, businesses must prioritize cost containment strategies because exceeding the budget is as equally unacceptable as breaching an SLA.

Debating between violating capacity limits, increasing costs, or impacting business-critical SLAs, becomes heated and sometimes leads to finger pointing. In the meantime, the customer experience suffers. The challenge becomes how to meet SLAs at the best possible price.

Solutions Overview (cont)

Monthly license charges can be the largest single monthly invoice for businesses that deploy mainframes. MLC are determined by a rolling four hour average (R4HA), which is a four-hour peak-capacity-usage during that month.

Both IT and business stakeholders desire a predictable and proactive way to meet SLAs, especially for critical workloads, with minimal impact to cost. Selecting the best possible combination of pricing options, monitoring variables, and resource management strategies is a business critical decision.

Intelligent Automation Improves Management of SLAs and Avoids Unplanned Spikes in Cost
CA Dynamic Capacity Intelligence is an easy-to-use, fully automated, and intelligent capacity management and real-time dynamic capping solution. The solution helps you to monitor and effectively manage SLAs, improving operational efficiency. With CA Dynamic Capacity Intelligence, you can achieve rapid value by enabling priority mainframe workloads at the required capacity to meet SLAs. Moreover, you can take the guesswork out of MLC pricing.

CA Dynamic Capacity Intelligence recognizes the difference between urgent, business critical tasks and lower-priority tasks that can be briefly postponed. The solution lowers the overall million service units per hour (MSU) during peak cost times by automating mainframe capacity balancing. Automated capacity balancing lowers MLC charges while simultaneously increasing the capacity for time-critical workloads.

Moreover, CA Dynamic Capacity Intelligence provides a single point of control to simplify and automate enterprise-wide deployment, monitoring, and reporting. The single point of control insulates staff from the underlying complexities of many aspects of their operations, including the overall spectrum of software portfolio usage and cost.

Critical Differentiators

CA Dynamic Capacity Intelligence is a complete solution that delivers the following critical differentiators:

- **Real capacity intelligence**. Unlike other optimization methods available today CA Dynamic Capacity Intelligence does more than simply move workloads to areas with more capacity. The solution maintains SLAs for critical workloads by continuously analyzing workload capacity use, and automatically moving capacity to where and when it is needed.
- **Flexibility**. CA Dynamic Capacity Intelligence offers the ability to work with, without, or within the IBM group capacity limit (GCL) for groups of LPARS and defined capacity (DC) for singular LPARs. You can prioritize MSU capacity based on workload priorities within an LPAR or across LPARs. Additionally you can define a flexible range for MSU use and automatically allocate the minimum usage. If defined minimum usages have already been exceeded, they can be defined to become the new baseline.
- **Power**. CA Dynamic Capacity Intelligence gives you the power to shift resources to deliver business innovation and to lower software costs across mainframe operations. With CA Dynamic Capacity Intelligence, you can balance the competing needs of customers, the IT organization, and the business. CA Dynamic Capacity Intelligence provide the business continuity that your CTO expects, delivers the mainframe platform ROI demanded by your CFO, and improves the overall customer experience.

Establish a Successful Foundation for Intelligent Mainframe Capacity Management

The ultimate goal is to empower IT to better manage service level agreements through automated, intelligent mainframe capacity management thus helping drive down MLC. But what are the top challenges that you might face when implementing SLA and cost controls?

Overcome Ten Common Challenges in Mainframe Capacity Management

How can CA Dynamic Capacity Intelligence help you conquer the challenges of mainframe capacity management? We will investigate ten common challenges and show how CA Dynamic Capacity Intelligence can help you conquer them.

Overcome Ten Common Challenges in Mainframe Capacity Management (cont.)

Understanding Your Monthly License Charge (MLC)

IBM charges enterprises for its proprietary software such as z/OS, DB2, CICS, IMS, WebSphere MQ, based on the highest usage over any given four-hour period (R4HA) during the monthly billing cycle. You can significantly lowered your MLC by using CA Dynamic Capacity Intelligence to allocate the available resources more efficiently. Furthermore, CA Dynamic Capacity can help you monitor, and therefore to understand, a number of other factors that impact your MLC. Where might your enterprise be overpaying for unneeded capacity? What can you do to address each of these impacting factors?

Figuring Out Your Real Peak Usage

Many enterprise customers do not realize that while an MLC billing period starts on the second day of a month, the billing period also includes the first day of the *next* month. This way of measuring peak usage, disadvantages many enterprises because their highest system usage occurs on the first working day of the month. The first day of the month is a high usage time because of the end-of-the-month processing for the previous month. The first day of the month is also, typically, the last day of the billing period.

Many enterprises strive to reduce the number of million service units (MSU) that they use during the month, only to be disappointed on the last day of the billing period. After experiencing this several times, they might give up on trying to economize. They, therefore, set caps too high and needlessly waste money on MSUs that they do not need.

What can you do to better predict and control your MSU use? CA Dynamic Capacity Intelligence can help you to define a flexible range for MSU usage. Even more important, the solution automatically allocates the minimum usage, modulates usage as needed, and ensures that usage never exceeds the maximum that you define. The truly intelligent solution can even help you plan for and avoid any negative impact from MSU-hungry events, such as end-of-the-month processing.

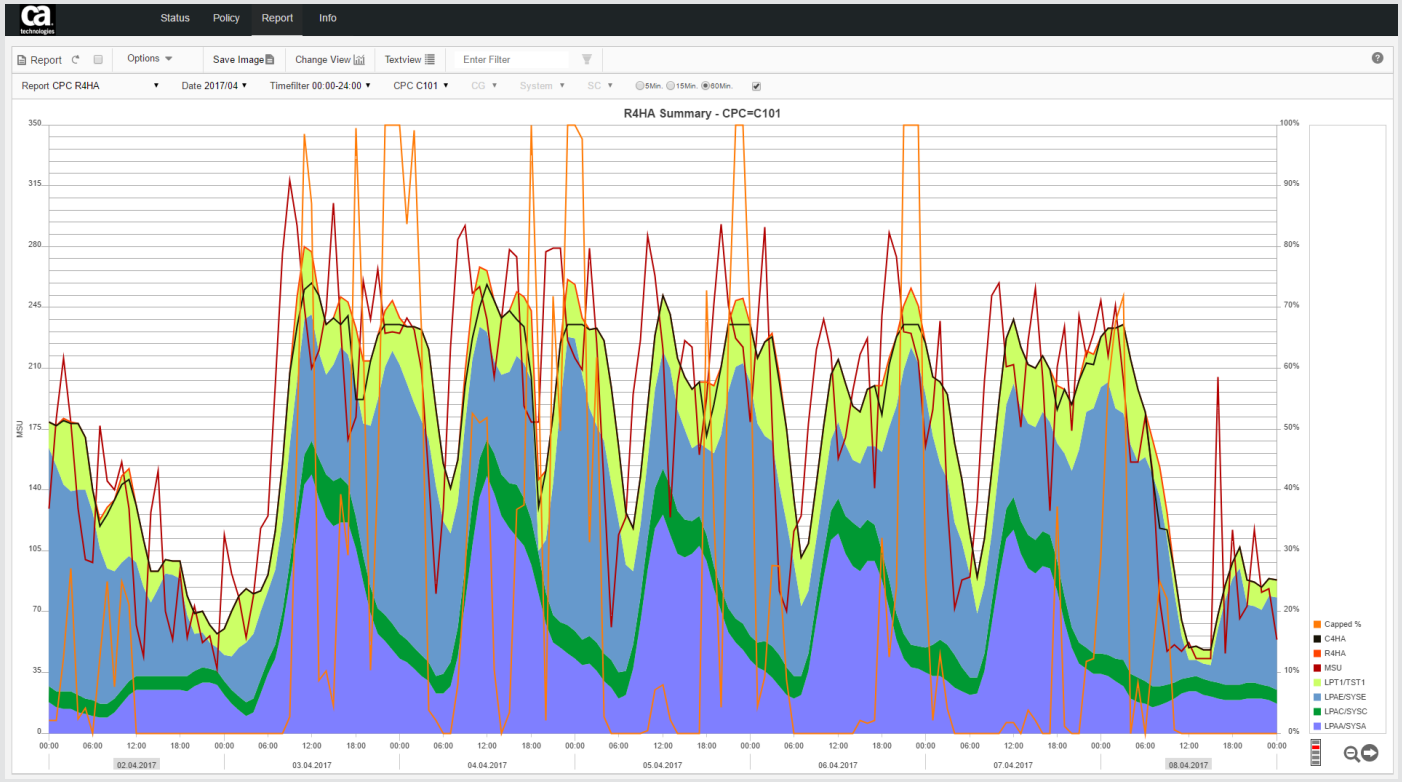

Easily view your peak real-time rolling four-hour average (R4HA), workload caps, and monthly license charge (MLC).

Realizing That GCL and DC Alone Are Inadequate

To control usage, IBM offers two metrics: defined capacity (DC) for single LPARs, and group capacity limit (GCL) for groups of LPARs. These mechanisms offer little flexibility because they need to be set and reset manually. Basically, DC and GCL only help to protect against very high usage. Whenever a GCL is reached, all MSU usage is divided based upon the previously defined settings. The problem is that these settings rarely reflect the real need during the actual peak time.

The only alternative is to define a flexible minimum and maximum MSU range that automatically allocates the minimum usages required to meet the workload. You should only add capacity when it is necessary to run the critical workload. Taking this approach ensures that the MSU usage over time is always as close to the required minimum as possible, while never exceeding the defined maximum. CA Dynamic Capacity Intelligence continuously monitors MSU usage and provides additional capacity where you need it without exceeding defined maximums.

Struggling to Balance Capacity Across LPARs Based on Workload Priorities

Manual capping is often based only upon the R4HA maximum for a single LPAR's DC, or a group of LPAR's GCL. However, the goal is to use the available capacity optimally. So, you want to know if one LPAR is struggling to execute a time-critical workload while, simultaneously, another LPAR is still running less urgent work. Furthermore, different LPARs may run similar workloads in identical workload classes but with a different priorities such as batch workload in production compared to a batch workload in development. The level of priority might even differ depending on the time of day, such as production batch workload during the day or during the night. Manual capacity management does not offer visibility into these complexities.

Using CA Dynamic Capacity Intelligence in conjunction with z/OS Workload Manager (WLM) gives you visibility into the complex, competing priorities for the available capacity. CA Dynamic Capacity Intelligence defines the levels of priority with a corresponding WLM Importance Level. It also enables you to define multiple capping policies. Each policy is activated depending on the day, date, and time of day. Policies let you rest assured that urgent work is always prioritized, while less crucial workloads are slowed down to avoid the need to add more capacity.

Paying for More R4HA Than You Actually Use

At many sites, the actual MSU peaks over a two-hour period in the early to mid morning. Because the R4HA reflects the usage of the past four hours, it is much slower than the actual MSU usage. In peak situations of less than four hours, the R4HA goes up while the MSU demand drops. The result is that enterprises are paying for capacity they do not use.

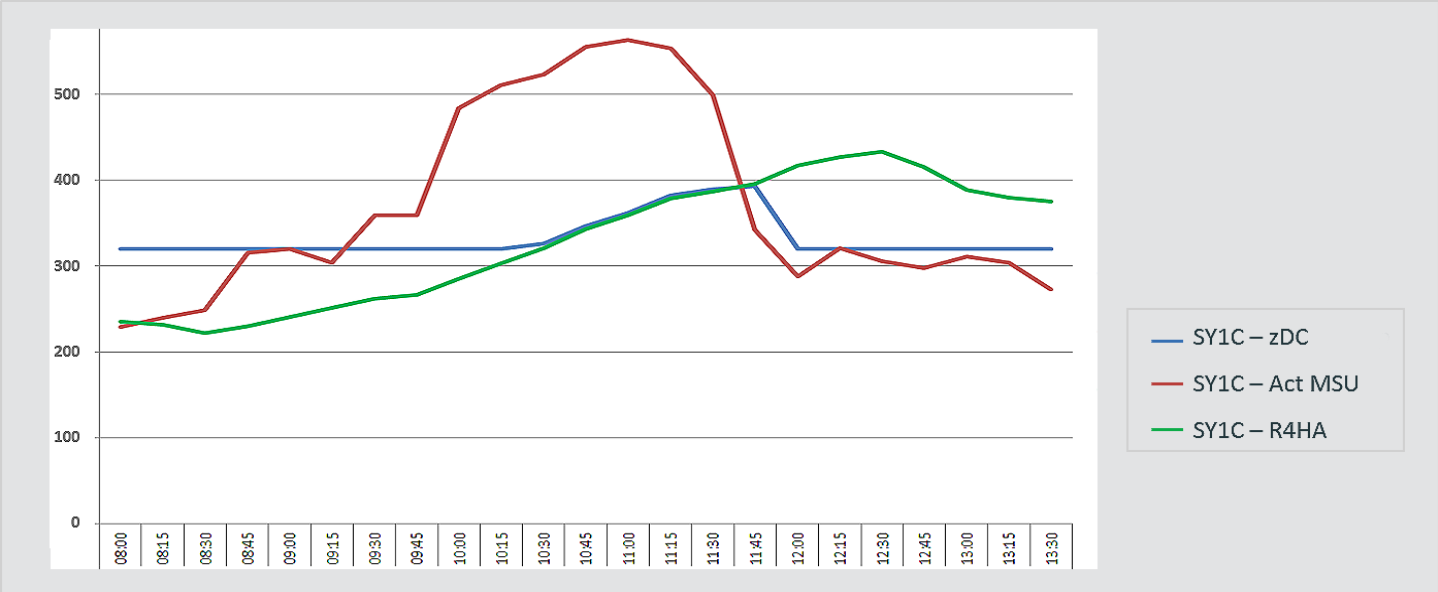

This figure shows the relationship between actual MSU, R4HA and the DC of the IBM z/os system.

Paying for More R4HA Than You Actually Use (cont.)
You can automatically perform capping without limitations with CA Dynamic Capacity Intelligence. If the current usage, actual MSU, drops below the R4HA, the defined capacity follows the current usage.

There are three advantages to this method:

- Users will not notice the capping as there is no limitation.
- The Sub Capacity Reporting Tool (SCRT) will take the lower of the two values of R4HA and DC.
- Capacity that is freed up may be used by other LPARs at no additional cost.

Capping Instead of Managing
Defined capacity (for single LPARs) and the group capacity limit (for groups of LPARs) limit MSU usage. They are simple ways to make sure that the MLC costs do not exceed the defined maximum. They are not MSU management tools!

Caps need to be recalculated regularly to guarantee that there is enough capacity to perform time-critical workloads. When you use caps too aggressively, do not monitor them closely, or use them in an environment where the system settings are not optimal, problems can occur during peak times. For example, online transactions may be slowed down while at the same time low priority batch workloads still run fine. For these reasons, many enterprises understandably refuse to cap their production workloads in any way.

However, when you use CA Dynamic Capacity Intelligence to recalculate your caps every five minutes and to monitor time-critical workloads, you guarantee that hight priority work gets done while you still deliver a high level of customer satisfaction.

Losing System Programmers and System Data to the Knowledge Drain
IBM z/OS systems programmers are quickly becoming a scarce resource that cannot easily be replaced. Each seasoned mainframe professional who retires or departs an enterprise leaves behind a major knowledge gap. When you use CA Dynamic Capacity Intelligence with the z/OS WLM feature, you put the MSU settings into intuitive software parameters. Moreover, you document these parameters so that the technical staff can simply monitor the performance of an automated program.

Using CA Dynamic Capacity Intelligence in this way provides two additional benefits:

- It frees up experienced staff to attend to more vital tasks.
- It facilitates training new staff.

Trying to Do the Impossible in Increasingly More Complex Environments
As more new platforms interlink with each other and use IBM z/OS systems as a major data hub and repository for massive data volumes, these systems become too complex and too difficult to manage manually. Manually allocating the optimal MSU capacity for a small environment is challenging; but doing it manually for a large environments is simply impossible.

Why try to do the impossible? Instead, let CA Dynamic Capacity Intelligence automatically calculate the optimal settings every five minutes. You can then focus on defining the parameters and rules that you want for each LPAR, depending on the LPAR specific workload.

Needlessly Increasing Cap Values
In many cases, system programmers raise the DC or GCL maximum values when demand threatens to overwhelm them. System programmers work under tight deadlines to avoid hitting maximum capacity. They often do not have time to gather the information about which LPARs to prioritize. They, therefore, must quickly make several educated guesses. How high should they raise the DC or GCL values, for which LPARs, and for how long?

CA Dynamic Capacity Intelligence, with automatic capping, takes the guesswork out of calculating cap increases because it automatically calculates cap increases every five minutes. This means that increases only have to cover usage needs for the next five minutes. These small MSU increases guarantee that time-critical work gets done as scheduled, using only the capacity that is needed. As soon as the actual usage drops below the R4HA line, CA Dynamic Capacity Intelligence recalculates capacity and rebalances workloads. As a result, DC value is reduced and some low-priority work is executed once again.

Needlessly Increasing Cap Values (cont.)
In the image below, the graph shows that the DC was manually raised from 490 to 550 at around 10:30 a.m. and set back to 490 at 4:30 p.m. The R4HA did not come close to the value of the DC, leaving unused capacity that could have been used for discretionary workload.

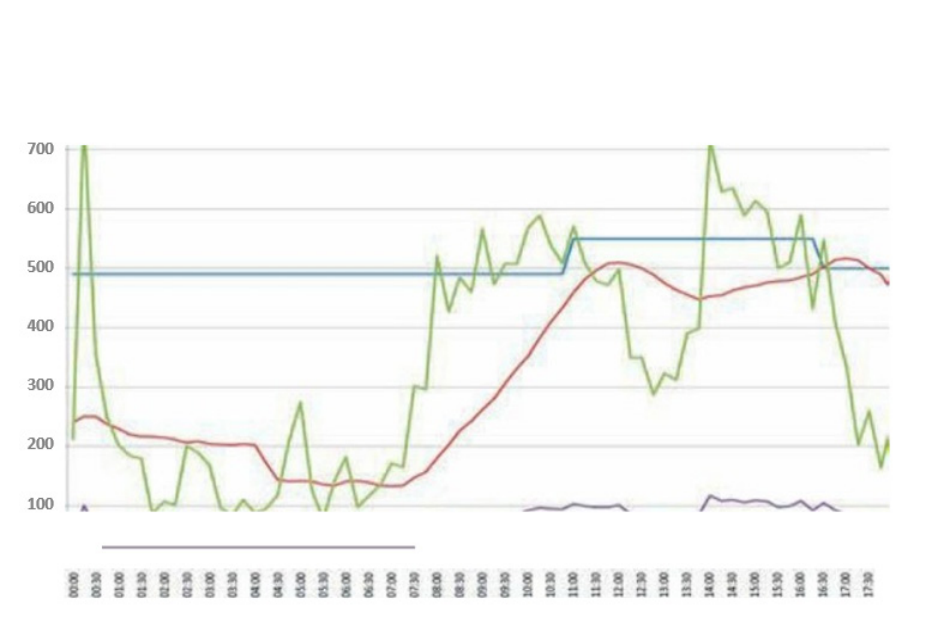

Manual DC management can result in higher caps and wasted capacity.

In the image below, the graph shows how automated capping in CA Dynamic Capacity Management works. With automated capping, the DC was raised to meet demand just before capping became necessary. At the same time, the lower priority workload was slowed down. Using CA Dynamic Capacity Management with automated capping resulted in a savings of around 70 MSU, all without impacting the critical workload. Assuming this enterprise runs IBM's z/OS, DB2, CICS, and NetView, the monetary savings are likely to be more than $400 per MSU, or in excess of $28,000 for this one month.

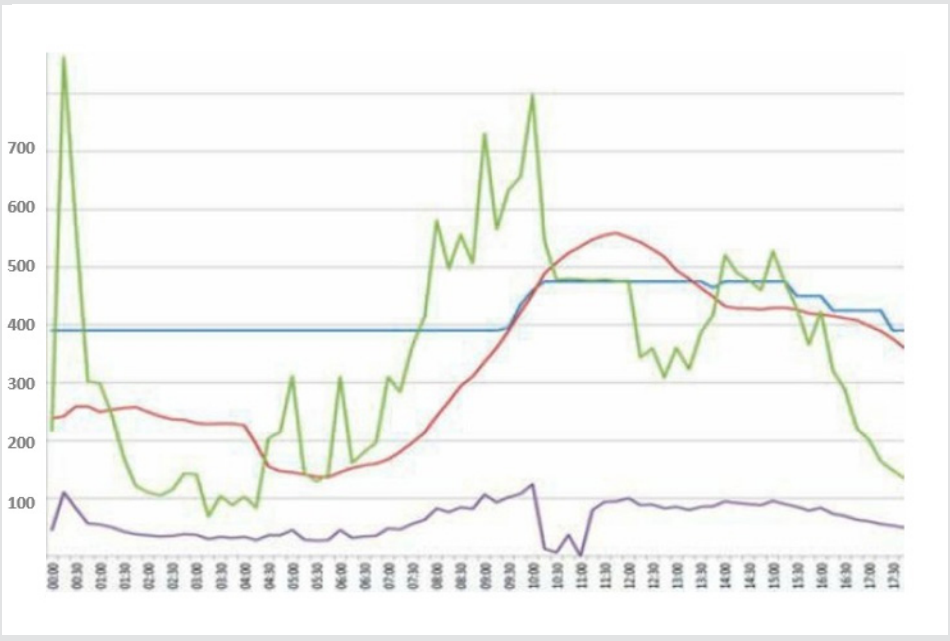

Automatic DC management can help reduce MSU usage much more effectively than manual DC management.

Making Decisions Without Real-Time Information

In many environments, development LPARs are limited in their MSU usage. People assume this is necessary. However, in many situations there is no, or almost no, financial gain made by capping. For example, during times that production LPARs do not need much capacity, some of that excess capacity could be given to development LPARs. Giving this capacity to development LPARs to use in application development can be done at no or little extra cost.

CA Dynamic Capacity Intelligence, through its LPAR-specific parameter settings, automatically provides real-time information about the financial impact of limiting or reallocating MSU usage.

Conclusion

CA Dynamic Capacity Intelligence provides proactive, predictable capacity management for optimizing system resources for prioritized workloads.

CA Dynamic Capacity Intelligence provides all of the following benefits:

- Ensures that priority workloads complete per SLAs with minimal impact to MLC.
- Dynamically moves available capacity to LPARs running critical workloads.
- Automatically balances capacity in real-time to fully utilize available capacity.
- Helps decrease human labor hours that are required for manual real-time capacity management.
- Provides complete transparency into capacity management.

To learn more about our mainframe capacity management solutions, and how they can help you achieve real savings in your environment, visit ca.com/dci.

CS200-SPK-0419 April 9, 2019