



COMPUTER
CORPORATION
OF AMERICA

Concepts
and
Facilities

MODEL 204

MANAGEMENT SUMMARY

TABLE OF CONTENTS

AUDIENCE	2
INTRODUCTION	3-4
The More Things Change, The More They Remain the Same	3
Meeting Current and Future Challenges	3
Mainframe Power on Any Desktop	4
Productive Application Development	4
Flexibility: Open to Change	4
MODEL 204 PRODUCT FAMILY OVERVIEW	5-9
Performance Options	5-6
<i>MP/204</i>	5
<i>Parallel Query Option</i>	5
<i>A Centralized View of Decentralized Operations</i>	6
Development Tools	6-7
<i>User Language</i>	6
<i>User Language/DB2</i>	7
<i>Workshop/204</i>	7
<i>WebGate Development</i>	7
Connectivity Products	8
<i>Horizon</i>	8
<i>Connect★</i>	8
<i>MQSeries</i>	8
<i>WebGate Connectivity</i>	8
Query and Reporting Tools	9
<i>Sage/204</i>	9
<i>Imagine</i>	9
<i>Access/204</i>	9
FEATURES, FUNCTIONS AND BENEFITS	9-16
Performance	9-11
<i>Advanced Indexing Technologies</i>	10
<i>Optimized Relational</i>	10
<i>Iterative Queries</i>	10
<i>Minimized Disk Activity</i>	11
<i>High Volume Update Facilities</i>	11
Productivity	11-12
<i>Rapid Iterative Development</i>	11
<i>SQL</i>	12
<i>End-User Productivity</i>	12
Scalability	12
Flexibility	12-14
<i>Dynamic Databases</i>	12
<i>Dynamic Binding</i>	12
<i>File Groups</i>	13
<i>Enhanced Performance and Capacity</i>	13
<i>Views</i>	14
<i>Complex Data Types</i>	14
Openness	14-15
<i>Adherence to Standards</i>	14
<i>SQL</i>	14
<i>OBDC and OLE DB</i>	15
<i>APPC/LU6.2</i>	15
<i>Portability</i>	15
<i>Connectivity</i>	15
Security	15-16
<i>Distributed Connectivity</i>	16

BUSINESS APPLICATIONS 16-18

Very Large Databases and Data Warehousing 16
Very High Volume OLTP and OLAP 17
Mainframe Data Servers 17
Distributed Applications 28
The Web 28

SUMMARY..... 19

THIS DOCUMENT EXAMINES:

- ▲ Concepts underlying the Model 204® database management system and related products
- ▲ Facilities offered by the Model 204 product family
- ▲ Applications of the technology

It describes how the flexible, open and scalable Model 204 product family powers some of the largest and most complex database applications around the world. You will also learn how Model 204 application development facilities improve productivity while making data easily available to users of most computing platforms.

AUDIENCE

This information is of most value to executives and managers who are new to Model 204 or who are considering introducing it into their organizations. It shows you how the product family can benefit your organization.

Another document, "Model 204 Concepts and Facilities: System Overview," addresses the needs of technical professionals who require a more detailed discussion of the Model 204 product family.

INTRODUCTION

Computer Corporation of America (CCA) is one of the top 30 privately held software companies in the world. Some of the largest corporate, government and non-profit organizations in the world depend on CCA for enterprise data management solutions that fulfill their most exacting information requirements.

Users of Model 204, CCA's mainframe database management system, often have exceptional demands: terabytes of data, millions of rows, thousands of concurrent users, open connectivity, secure environments, and almost instant response to complex queries. While these customers need to manage very large databases effectively, the solution must also economically satisfy their smaller departmental information requirements.

For more than thirty years, CCA has fulfilled its customers' needs for scalable, high-performance, flexible, robust, and secure data management.

THE MORE THINGS CHANGE, THE MORE THEY REMAIN THE SAME

While much has changed since the founding of CCA as a bespoke software development company in 1965, some things remain the same:

- ▲ No matter how rapidly computer processing and storage price/performance ratios improve, it is never enough. The more data and computing resources that organizations have, the more valuable uses they find for it.
- ▲ There are never enough programmer and analyst hours. Software developers must be exceptionally productive to tackle the application backlog at a reasonable cost.
- ▲ No technology should be an island. The greatest gains come from seamlessly integrating technologies across the enterprise.
- ▲ Tomorrow will not be like today. Your database management technology must easily adapt to constantly changing requirements.
- ▲ In the end, what matters most is quick, easy, and secure access to the information people need to maximize value for their enterprises.

These requirements drove the development of CCA's innovative technologies. They remain the principles underlying the advancement of the Model 204 family of data management products today. The scale and flexibility of Model 204 enable the technology to deliver rapid solutions for high-end OLTP and OLAP applications, while providing the most cost-effective solutions for the smaller user.

MEETING CURRENT AND FUTURE CHALLENGES

While years of practical use prove its underlying principles, Model 204 continues to evolve to meet the changing needs of business and incorporate technology advances. Every new release of Model 204 contains significant enhancements to:

- ▲ Take best advantage of new technologies
- ▲ Improve performance and scalability
- ▲ Provide open connectivity
- ▲ Protect and leverage our customers' investment in information technology

The advancement of the Model 204 product family is a road that never ends. As the technology and business environments change, CCA is committed to making the necessary investment to ensure that we continue to meet our customers' evolving information management requirements and preserve their investment in Model 204 data and applications.

MAINFRAME POWER ON ANY DESKTOP

Because it is tuned for mainframe platforms – OS/390, VM, and VSE – Model 204 takes optimal advantage of the underlying operating system. Consequently, it can manage databases with terabytes of data, hundreds of millions of records, and thousands of concurrent users.

Furthermore, because of its unique indexing technologies and other advanced techniques, Model 204 performance is immune to the exponential degradation in performance experienced with other technologies, even as databases grow very large.

Even as PCs replace mainframe terminals on users' desktops, users still retain mainframe performance. The Model 204 product family offers seamless access to data from desktop systems, whether it be a PC, UNIX workstation, or a Web browser running on any platform.

PRODUCTIVE APPLICATION DEVELOPMENT

Organizations face an ongoing shortage of skilled application developers. Providing tools to make developers as productive as possible is an essential element of every information technology manager's job.

User Language™ provides an integrated fourth generation environment for building Model 204 applications. CCA customers find that User Language can achieve a greater than ten-to-one improvement in programmer productivity over third-generation languages.

Workshop/204™ enhances application development even further by integrating User Language with an interactive and efficient environment for performing all phases of development, from file design to testing.

Furthermore, because Model 204 is accessible through industry standard SQL, you can also use existing SQL skills to quickly develop Model 204 applications.

FLEXIBILITY: OPEN TO CHANGE

Because your business environment never stops changing, Model 204 is open to change. Database administrators can make many database structure changes without taking systems off-line. Consequently, your databases evolve in step with your business, without the need to stop or even slow operations.

The remainder of this document examines the features and benefits of the Model 204 product family in more detail.

The **Model 204 Product Family Overview** section describes the members of the product family and how they work together to benefit your organization.

In the **Features, Functions, and Benefits** section you will learn about the unique facilities of Model 204 and its related products. This section focuses particularly on the value you can receive from these facilities.

The **Business Applications** section examines specific business requirements and how you can apply the Model 204 data management family of products to meet them.

MODEL 204 PRODUCT FAMILY OVERVIEW

More than 1,000 organizations around the world benefit from Computer Corporation of America's Model 204. Its flexibility, scalability, and high performance have enabled customers to grow from small business operations to enormous enterprises, rapidly and cost-effectively.

These customers depend on the performance, openness, flexibility, scalability, and productivity of Model 204 to drive their most critical business applications. The Features, Functions and Benefits section beginning on page 9 discusses how Model 204 achieves these qualities.

However, Model 204 is only one important component of the product family. This section describes some of the other components and how they can enhance your business.

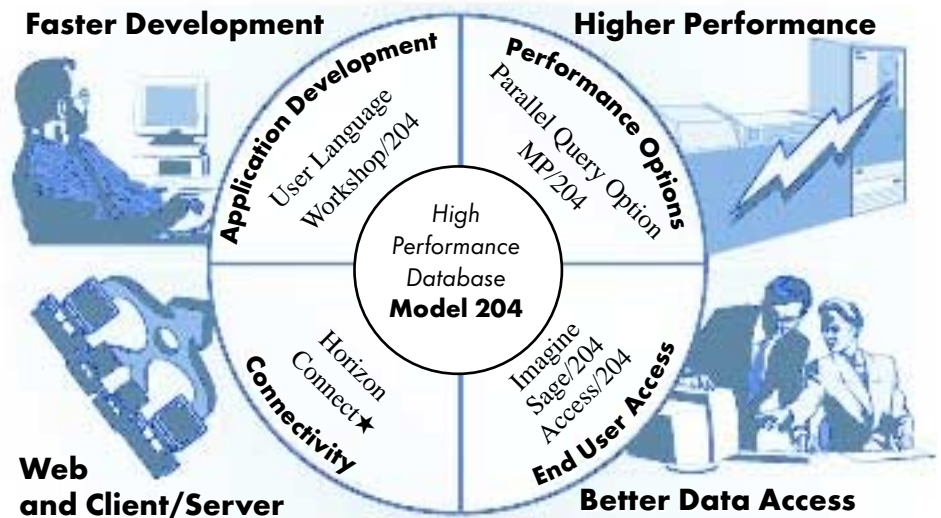


Figure 1: Model 204 Product Family

PERFORMANCE OPTIONS

While most organizations find the inherent power of Model 204 more than adequate to meet their needs, two options, Multiprocessor/204 (MP/204) and Parallel Query Option (PQO), can boost its performance even further.

MP/204

MP/204 provides users of IBM and compatible multi-processor mainframes with a way to balance the application workload among processors by off-loading between 40 and 90 percent of application processing to additional processors. Its sophisticated workload balancing increases overall system throughput and enables support for additional users without the need for any additional equipment.

MP/204 provides the means to scale up your systems and data – more users, more data, more functionality – without upgrading your hardware. With MP/204, your systems can grow economically to keep pace with your business.

PARALLEL QUERY OPTION

Another scalability option is PQO, which distributes a single query across multiple Model 204 databases.

The distributed query is transparent to users. To them, it appears to be running against one seamless database. PQO parcels the query out to all participating systems, which process it against their local databases. PQO then consolidates the results and returns them to the user.

PQO provides almost unlimited expansion possibilities. You can connect it to as many as 2,040 instances of Model 204 and process a single query against as many as 256 Model 204 instances in parallel. PQO can, thus, significantly reduce query response time, particularly for complex queries against very large databases.

A CENTRALIZED VIEW OF DECENTRALIZED OPERATIONS

PQO is more than just a performance option. It is also an effective way for head office managers and executives to gain an enterprise-wide view of decentralized operations. Even if all subsidiaries maintain their own databases, a single PQO query can provide management with consolidated corporate information.

DEVELOPMENT TOOLS

Application development efficiency is a concern for every enterprise. Model 204 customers have an advantage. They can use some of the most productive development tools available, such as User Language and Workshop/204.

USER LANGUAGE

User Language offers a powerful fourth-generation environment for building Model 204 applications. It integrates data access and business process logic into a single intuitive programming language that significantly improves development productivity.

User Language ensures application flexibility by accessing table layouts, column names, string lengths and other data implementation details only when needed. Consequently, applications automatically recognize dynamic changes to the database structure without the need to modify or recompile programs. Furthermore, unique Model 204 technology delivers the very high performance that eludes traditional dynamic binding schemes.

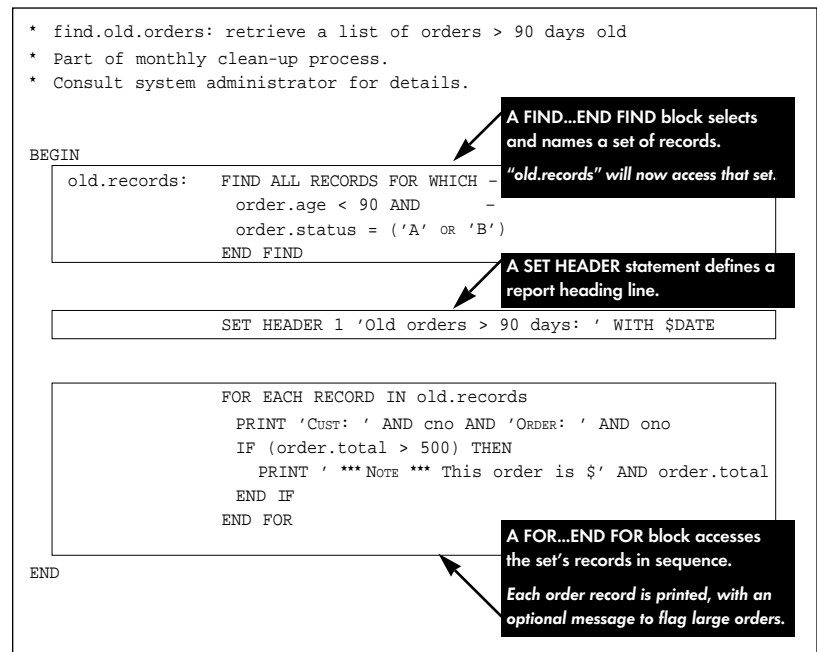


Figure 2: Sample User Language Procedure

User Language achieves its productivity gains through a number of techniques:

- ▲ High-level statements and functions perform work that would require several statements in other programming languages. This reduces the programming effort and makes applications easier to maintain.
- ▲ Many statements create or manipulate sets of rows. Users can reference sets by meaningful names, such as "old.records" in the example in Figure 2. Working with sets as a single entity, rather than working with one row at a time, greatly simplifies many operations.
- ▲ Because User Language accepts long names (up to 512 characters), uses meaningful keywords, provides data dictionary services, and allows in-line comments, it is self-documenting. Thus, it eliminates considerable external documentation and the effort required to create it.
- ▲ Experienced users can make use of time-saving abbreviations, such as "FD" instead of "FIND ALL RECORDS FOR WHICH."

User Language programs can be run on-line or as batch processes, giving:

- ▲ Immediate answers to urgent questions by composing, editing, and running User Language requests interactively on an ad hoc basis
- ▲ Minimal impact to online production systems by deferring requests to batch execution
- ▲ Significant time savings by reusing a request later in a session
- ▲ Improved productivity and shared User Language knowledge by saving requests in procedure libraries for subsequent on-line or batch invocation

In addition, User Language is fully integrated with Model 204, thereby vastly improving its performance over other fourth-generation languages.

USER LANGUAGE/DB2

User Language/DB2 provides connectivity between DB2 data and Model 204. It harnesses the powerful fourth-generation capabilities of User Language to provide significant productivity gains when accessing or updating DB2 information.

Like Model 204 User Language, User Language/DB2 does not sacrifice performance to gain its outstanding ease-of-use. Instead, it integrates the productivity of a fourth-generation language, the performance and power of traditional third-generation languages, and the data independence and flexibility of relational data manipulation to take best advantage of your DB2 data.

By providing the ability to marry Model 204, DB2 and other data, User Language/DB2 leverages the return on your data investment by integrating heterogeneous, enterprise-wide databases.

WORKSHOP/204

Workshop/204 provides a complete development environment, integrating User Language with a screen and report painter through a single developer interface.

Several time-saving facilities help developers modify applications quickly and easily:

- ▲ The procedure editor lets programmers quickly define, edit, and delete User Language procedures without using complex syntax.
- ▲ The query/update facility lets developers quickly create and modify database views without programming.
- ▲ The data dictionary automatically documents every aspect of the database and reports the impact of database changes on individual applications. Thus, programmers can perform routine maintenance confident that their changes will not affect other applications.

WEBGATE DEVELOPMENT

More than just connectivity software, WebGate™ provides a powerful, wizard-driven development environment for the rapid development and deployment of Internet applications, which include Model 204 data and transactions.

CONNECTIVITY PRODUCTS

HORIZON

Most organizations use a variety of hardware and software platforms. Horizon's open computing architecture allows these diverse systems to work together seamlessly.

By integrating an industry standard communications protocol with CCA's powerful User Language, Horizon™ allows you to manage complex interprocess coordination and communication easily. As a result, you can create seamless cooperative application environments across address spaces on one machine, or across multiple computers running a variety of operating systems – from LAN PCs to platforms such as SUN, HP, DEC, and IBM. Applications simply initiate a conversation with the system managing the required data.

In an integrated Horizon environment, information is location-independent. Users and applications request information without regard to its physical location. Horizon then directs the request to the appropriate system. You can, therefore, easily develop applications on one machine and install them elsewhere.

And, as your systems grow, you can logically group access to data from new locations using symbolic definitions, thus simplifying the process. Furthermore, you can add or modify definitions without interrupting the production environment.

In addition, Horizon offers facilities for managing the distributed environment, including comprehensive and flexible security and 24-hour network operations.



Figure 3: Seamless Mainframe-Desktop Connections

CONNECT★

Connect★™ provides connectivity between Model 204 databases and desktop Windows and UNIX systems.

Connect★ allows the user to take full advantage of client/server and Internet architectures. On the client side, end users can employ the PC and UNIX software already familiar to them. On the server side is the full power of Model 204, ensuring high-performance database management. Connect★ marries the two through industry-standard SQL. Alternatively, with CCA's innovative use of the ANSI standard SQLExecDirect facility, desktop application programmers can access Model 204 data using the full power of CCA's User Language coupled with the data analysis and graphical presentation facilities of industry standard PC software.

MQSERIES

MQSeries is rapidly emerging as a standard for industrial strength, cross-platform connectivity. MQ/204 adds Model 204 parameters and commands and User Language syntax and functions to enable Model 204 applications directly to benefit from this leading edge product.

MQ/204 delivers timely and flexible capabilities for rapid integration of Model 204 within our customers' Message Queuing Middleware solutions.

WEBGATE CONNECTIVITY

On the World Wide Web, WebGate provides a powerful link between Windows NT Web servers and the mainframe; allowing rapid deployment of highly scalable intranet and extranet applications linking Model 204 enterprise systems to today's rapidly expanding, nomadic community of Internet users.

QUERY AND REPORTING TOOLS

SAGE/204

Sage/204™ provides a transparent link from Model 204 to SAS. Furthermore, it employs syntax familiar to both Model 204 and SAS users, thus reducing learning costs and improving productivity.

Using Sage/204, SAS programmers can write data extracts with only limited knowledge of Model 204 language constructs. To produce a SAS report or graph, they simply code a PROC SAS statement following the same syntax rules as any other SAS procedure. They can then enhance the results by using other SAS statements in conjunction with the PROC SAS statement.

Likewise, Model 204 programmers can use the constructs familiar to them to create SAS extracts.

IMAGINE

Imagine® and its Windows-based graphical user interface companion, Imagine/CUA®, are end-user query and reporting tools that provide seamless access to a wide range of IBM data file types. Imagine and Imagine/CUA enable users to combine data from Model 204, DB2 for OS/390, DB2 for VSE & VM, IMS, DL/1 and VSAM and other sources into a single, easy-to-understand report or query.

Imagine handles a variety of business requirements, from regularly scheduled batch production reports with fully specified formats and layouts, to ad-hoc on-line queries with system defaults for format and layout. Imagine is available for OS/390 and VSE operating systems.

ACCESS/204

Access/204™ is a query and reporting tool that users employ to access data by simply entering request criteria into an online form. It allows them to access information by views, which may be a selection of tables, columns and rows or a combination of tables that presents the user's perspective of data, rather than its physical layout.

In addition, a view may provide computed fields that are not stored in the database, but rather derived from other fields. For example, a user view may include the option to convert dates of birth to age ranges for ease of understanding and analyzing this data, rather than having to require Information Services to maintain this volatile information in the database itself.

FEATURES, FUNCTIONS, AND BENEFITS

PERFORMANCE

Why is Model 204 the best performing and most cost-effective database for real-world applications? Because Model 204 represents the culmination of many years of research and development specifically targeted at solving the real-world needs of large-scale customers.

For many years, before Model 204 came into being, CCA consultants were developing, honing, and refining data management techniques to resolve customers' needs where other, packaged technologies proved inadequate. Model 204 brings together advanced techniques for rapid data management, which other database vendors are still today striving to add on to their software. With Model 204 CCA not only continues to develop and leverage new techniques and technology, but we continue to enhance the performance of Model 204's existing capabilities. With the last four releases of Model 204 CCA has increased throughput by almost 300%, yet at the same time decreasing CPU consumption by 34% and physical I/O by 39%. This has enabled our customers to grow their business by several orders of magnitude without the necessity of expensive hardware upgrades.

A number of features contribute to the well-deserved reputation of Model 204 as the highest performing, most cost-effective database available:

ADVANCED INDEXING TECHNOLOGIES

CCA invented bit-mapped indexing and has continued to enhance the efficiency of this technology. It is only in recent years that some other database vendors have come to recognize bit-map indexing as the fastest form of indexing for many types of queries. Other products have introduced their own versions of bit-map indexing without the benefit of the many years of real-world experience that CCA continues to use to enhance the speed and efficiency of this technology, both for OLAP and OLTP applications.

Bit-mapping is the best technology for many types of queries and transactions. CCA’s experience has shown that, at times, other indexing techniques are superior. Therefore, Model 204 is not limited to bit-map indexing. It automatically combines the use of bit-mapping, hashing, B-tree, and record list technologies to optimize speed and efficiency of database access.

OPTIMIZED RELATIONAL MODEL

The relational model specifies how data structures should appear to the user. However, it does not define the physical implementation of the logical structure. Model 204 optimizes the relational model by, among other techniques, incorporating “embedded tables.”

One of the precepts of the relational model is that tables should not contain “repeating fields.” Employee skill is an example of a repeating field since employees can have multiple skills and the number of skills varies for each employee.

The relational model requires that repeating fields reside in a separate table, linked to the primary table through a common value, such as the employee number. However, by separating values into physically distinct tables, the time to access both the primary and subordinate data can grow dramatically. If you must access all the data most of the time, this can significantly

slow your applications. Embedded tables eliminate the problem while remaining true to the relational model.

To the user, an embedded table looks like a separate table, just as the relational model says it must. However, at the physical level, Model 204 stores the repeating values tightly coupled to the related primary data. Thus, Model 204 retrieves both the primary and subordinate data with a single database I/O. This significantly improves performance for applications that usually access all the data together.

Employee Table

ID	Name	Skills Table	Address

Figure 4: Embedded Tables

ITERATIVE QUERIES

Model 204 offers an iterative query facility that provides an efficient way for users to refine searches and data selections.

For example, an office supply vendor might first ask, “Which customers ordered office furniture last year?” He or she might then ask, “Of those, how many bought the top-of-the line models?” Using its set processing capabilities, Model 204 searches only the results of the first query to find the answer to the second. This can eliminate considerable processing, particularly if the database is very large.

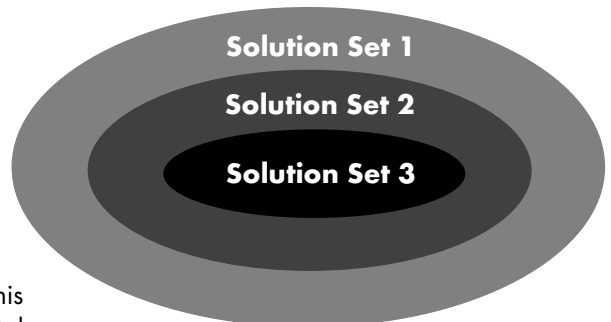


Figure 5: Iterative Queries

Iterative refinement of queries can continue until users find the precise information they need.

MINIMIZED DISK ACTIVITY

Moving data to and from disk drives is the primary bottleneck for most data intensive applications. Model 204 employs aggressive indexing, caching, and physical data storage strategies to minimize the amount of disk I/O and, thereby, eliminate the constraints.

HIGH-VOLUME UPDATE FACILITIES

Large organizations require robust, industrial strength databases capable of managing exceptionally high update volumes very efficiently. Model 204 offers several facilities to ensure that it can accommodate the most demanding of requirements, including:

- ▲ Model 204 locks database resources at the most granular level to limit the possibility of data contention.
- ▲ Model 204 can automatically lock and release rows as required, further reducing possible contentions.
- ▲ The file group facility allows you to partition data and perform high-volume update operations against the partitions in parallel (see File Groups on page 13).
- ▲ For instances when indices do not need to be locked during updating, Model 204 allows you to defer updating the indices until after it completes all the changes to the data. It then automatically sorts the index updates in the most efficient order for processing and applies them in bulk, thus greatly reducing the time required to update large volumes of data.

PRODUCTIVITY

Productivity includes a number of aspects. As the above section showed, Model 204 provides exceptional performance to ensure that the users can complete their jobs quickly. However, productivity also refers to the speed of application development and the brevity of required training. This section discusses general Model 204 productivity issues.

RAPID ITERATIVE DEVELOPMENT

User Language and Workshop/204 (see pages 6 and 7) allow organizations to compress the application development life cycle significantly and improve application quality, using a process known as rapid iterative development.

This process, with the supporting tools, integrates all aspects of the application life cycle – prototyping, development, administration, maintenance, and operations – in a single, developer-oriented, on-line environment. Integrated database, screen, report, communications, and support services simplify the process. Consequently, developers require only months to deliver projects that would take years using traditional techniques.

Organizations report that common bottlenecks in the design, coding, and testing phases disappear when using rapid iterative development. Furthermore, prototyping delivers results to end-users much sooner so that you can resolve problems before they become entrenched in the applications. As a result, you can deliver higher quality solutions faster.

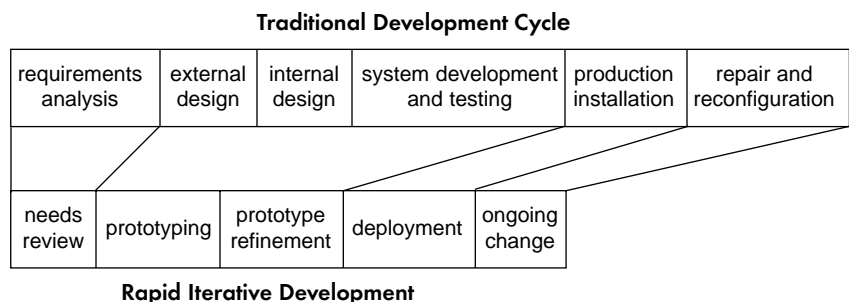


Figure 6: Rapid Iterative Development speeds implementation, yet improves results

SQL

SQL is the industry standard data access language. Therefore, there is already a considerable pool of SQL skills and experience in the workforce. Because Model 204 supports standard SQL access, new application developers can become productive almost immediately.

SQL alone may not satisfy the needs of business applications with very high transaction loads or complex data manipulation. Model 204's SQL implementation therefore exploits the ANSI standard to extend the SQL standard. It lets the developer take advantage of the full power of Model 204's integrated 4GL, by invoking User Language stored procedures directly from the desktop as a seamless component of a client/server or Web-based application.

END-USER PRODUCTIVITY

Productivity also extends to end-users.

- ▲ Powerful query and reporting tools such as Imagine and Access/204 (see page 11) put Model 204 data at their fingertips.
- ▲ Sage/204 (see page 11) allows them to combine the performance of Model 204 with the advanced reporting and statistical functionality of SAS.
- ▲ Connect★ (see page 11) transparently delivers Model 204 data directly to their familiar desktop applications, such as Microsoft Access, Excel, and Lotus 1-2-3.

SCALEABILITY

Business, government, and military organizations around the globe use Model 204 databases to store hundreds of millions of records and terabytes of data. Many of these customers support hundreds or even thousands of active users.

In addition, MP/204 and PQO (see page 7) can extend the power of Model 204 even further, making it suitable for the largest of databases and the most exacting of application environments.

FLEXIBILITY

Model 204 is also one of the most flexible mainframe database management systems in the world. As outlined below, its features allow your databases to keep pace with your business.

DYNAMIC DATABASES

It is a cliché, but change is constant. You cannot afford to stop your business to adapt to ever-present change. Model 204 gives you the flexibility you need to keep pace. At any time, without stopping your systems, it allows you to dynamically:

- ▲ Define new files, tables, columns, indices, and relationships between tables.
- ▲ Relax or tighten security.
- ▲ Change column attributes and usage conditions.
- ▲ Add, amend, and delete user profiles.

DYNAMIC BINDING

Dynamic binding gives User Language applications the ability to adapt to many database changes, without modification or even recompilation. Simply, dynamic binding means that User Language procedures do not determine the structure of the database – table layouts, column names, string lengths, and other structural information – until the procedure runs.

Other programming languages typically bind this information to programs at the time the developer compiles the program. If there is a requirement to, for example, enlarge a column due, perhaps to Year 2000 compliance, developers must recompile the program and reload it in the production environment. However, with Model 204, you merely change the database structure on-line and programs automatically recognize the new structure. Those that do not use the column, but nevertheless access the table, will not require any reprogramming or recompiling.

In traditional environments, dynamic binding often leads to performance problems. Therefore, despite its benefits to application developers and database administrators, dynamic binding is rarely available for production environments. However, Model 204 is not a traditional environment. Its database technology, combined with the sophistication of User Language make dynamic binding efficient and practical. Column enlargement and adding columns can be accomplished without the need to reorganize the database.

FILE GROUPS

File groups (also known as 'partitioned tables'), is yet another concept pioneered by CCA in its work to solve real-world problems with innovative technology. As the name suggests, a file group is a means of referring to a group of files as a single logical entity. A file is the physical Model 204 structure that contains all or part of one or more Model 204 tables.

While some other databases offer a similar concept that allows users to refer to a union of tables, this should not be confused with the file group facility of Model 204. The union facility of other databases generally requires that the component tables be identical or nearly identical. A Model 204 file group, on the other hand, can contain files that have the same structure, but they can also contain files with just a generally similar structure, or even files that share just a single common column.

File groups can improve productivity significantly when creating applications. All commands and User Language statements that refer to a file group automatically perform operations on each member of the group in turn (subject to normal security constraints).

Furthermore, file groups can be either persistent or dynamic – even being changed programmatically to suit the requirements of each user of each application. When you change the components of the file group, applications immediately and automatically work with the newly constituted group, without modifying or recompiling the application.

ENHANCED PERFORMANCE AND CAPACITY

Using file groups to partition data into separate files, while allowing users to view them as a single entity, also provides opportunities to enhance performance and capacity:

- ▲ When loading large volumes of data, you can load or update all files in a group in parallel, thereby significantly reducing the total time required for the operation.
- ▲ You can use Parallel Query Option (see page 5) to distribute file group queries across multiple systems, thus providing even greater on-line transaction processing (OLTP) and on-line analytical processing (OLAP) performance.
- ▲ Splitting very large tables across files in a file group provides virtually limitless capacity. Using this technique, Model 204 supports customers who have databases with logical table sizes numbering in hundreds of millions of rows. File group technology will enable them to grow or create multi-billion row tables and still remain within the capacity of Model 204 to manage their data efficiently.
- ▲ File groups also provide more efficient access to very large databases and time series data. For example, one Model 204 customer has a single table of approximately 300 million records partitioned in a file group. Each month, the customer adds a file for each financial ledger to the file group – about 250 files over a two year period. As users run queries against the file group, the software dynamically re-groups the database so that they deal with the smallest possible physical entity that will satisfy their request.

VIEWS

Different users have different perspectives on data. Yet, each user may look at data in the same way day after day. If, for instance, a user always looks at Table X joined to Table Y, specifying that join every time can be cumbersome. Likewise, a particular user may not be interested in all columns and rows of a table. The extra data only gets in the way.

A view is a way to present users with just the information they need, in precisely the structure they need it. Database administrators can define views that predefine relationships among tables, and filter out certain columns and rows. Users are then spared the need to perform those data definitions.

A view can also contain a derived field that does not physically reside in the database. For example, metric weight may be derived from the American measure. Derived fields spare users the work of performing the calculations and minimize the risk of human error.

A view is a logical interpretation of data rather than a physical definition of its structure. Therefore, if users' requirements change, you can change the views without restructuring the underlying data.

Views are also one component of Model 204 security. Database administrators can limit some users to accessing data only through views that exclude sensitive data.

COMPLEX DATA TYPES

Today, data is often more than just words and numbers. It can be maps, pictures, signatures, video, sound, and other nontraditional information. Model 204 facilitates these diverse data requirements through the support of graphical images, variable-length strings, time-series, and other complex data types.

OPENNESS

Openness means different things to different people. To some, it is the adherence to standards so that they are neither tied to proprietary technologies, nor need to acquire specialized skills. To others, openness means portability – the ability to take an application and move it, unaltered, to another computing platform should that prove to be advantageous. Still others define openness in terms of connectivity – the ability to seamlessly and transparently share data among different platforms.

In reality, openness includes all of these criteria. It is the provision of business benefits through a solution that can easily extend across the enterprise and beyond, both today and into the future, without regard to the underlying technologies. As described below, Model 204 delivers all the qualities of a truly open solution.

ADHERENCE TO STANDARDS

SQL

Developers can access Model 204 data using SQL that is fully compliant with both the SQL Access Group Call-level Interface and Federal Information Processing Standard Publication 127.1.

To make applications more portable, and eliminate the need for additional developer training, the SQL Data Manipulation Language (DML) supported by Model 204 is pure industry standard SQL. Integration of Model 204 stored procedures is accomplished using the ANSI standard SQLExecDirect facility.

The Model 204 SQL Data Definition Language (DDL) includes one extension to the industry standard – the ability to define embedded tables. This allows database administrators to optimize the relational model at the physical level. Despite this DDL extension, at the logical level, developers and end-users can still access data as fully normalized relational tables.

Because Model 204 includes SQL support in its nucleus, rather than as an external translator, SQL-based access to Model 204 data is as efficient as access through User Language.

ODBC AND OLE DB

Open Data Base Connectivity is a *de facto* standard for database access from the desktop. Object Linking and Embedding for Data Bases is Microsoft's strategic direction. Model 204 supports both ODBC and OLE DB access through Connect★.

APPC/LU6.2

To ensure open connectivity to other computing platforms, Horizon fully supports Advanced Program-to-Program Communications (APPC), also known as LU6.2. APPC is an IBM standard that provides a means by which different types of computers, from a number of vendors, can communicate with each other.

PORTABILITY

Model 204 applications are portable across all IBM and compatible System/370 and System/390 platforms. Consequently, you can easily migrate programs from departmental systems to the largest of mainframe computers. Furthermore, should you need to change mainframe operating systems, applications can move, unmodified, between OS/390, VM, and VSE. The result is the ability to scale up your systems to accommodate business growth, without jeopardizing existing application investments.

With Model 204's ability to easily connect to other platforms, including Web browsers, client applications can achieve unprecedented portability while still receiving the performance and security benefits of a Model 204 server. Because Web applications are platform-independent, you can create Web-enabled Model 204 applications that can run against any server.

In addition, because of the open nature of the Web, these applications can access data from any connected server, combining information from Model 204 with information from other servers. This any-to-any capability means that you can easily integrate Model 204 databases into open information architectures.

CONNECTIVITY

Connectivity is another Model 204 strength. Horizon, MQ/204, and Connect★ provide seamless and transparent connections between Model 204 and most hardware, operating system, and database management system platforms. Data can reside anywhere in your organization and applications can access it without knowledge of the supporting technology or even its physical location.

SECURITY

Protecting the integrity and confidentiality of data is important to all organizations. Model 204 offers extensive security features to protect your enterprise information. Furthermore, because you can restrict access by column, row, file, procedure, terminal, and other elements, the granularity of protection can be as broad or fine as you require.

In addition, Model 204 provides interfaces that support popular third-party security software, such as ACF2, RACF, and CA-TopSecret. These interfaces allow you to define common security schemes for all your data, even if you use a variety of data management technologies.

DISTRIBUTED SECURITY

Today's distributed computing environments compound the security issues. Instead of a single system managing security for your network, now you must address security on many systems, both clients and servers. Horizon responds to the need for distributed security using a "trusted partners" protocol that recognizes three security levels:

- ▲ **Trusted Requests** include the user name and Already Verified Information (AVI).
- ▲ **Nontrusted Requests** send only the user name and password.
- ▲ Requests with no security contain no user information.

Client requests must match the trust level of the Model 204 server. To provide flexibility and control, you decide the access level of each server. Horizon then distinguishes between trusted and nontrusted partners.

To enhance your security environment, Horizon also allows you to monitor network definitions for usage, as well as dynamically modify, add, and delete security definitions.

BUSINESS APPLICATIONS

CA customers find that Model 204 serves a variety of business requirements. The inherent qualities of the database management system and supporting products make it the preferred solution for many of the exacting applications that businesses and government organizations seek to implement today. The following sections present an overview of a few types of applications in which Model 204 excels.

VERY LARGE DATABASES AND DATA WAREHOUSING

Retailers retain longer and more detailed histories of customer interactions so that they can serve customers better and, thereby, earn more business. Meteorologists collect more precise climatic data and analyze it over longer periods so that they can better predict the weather and warn people of impending disasters in time to take appropriate action. Manufacturers collect minute details about manufacturing processes, finished goods, assemblies, sub-assemblies, and raw materials so that they can better analyze and eliminate defects.

These collections of historical data, often referred to as data warehouses, can grow to be enormous. Similarly, operational databases have also grown over the years as their use expands to cover more business activities both internally and in dealings with customers, suppliers, business partners, and governments. Today, information technology managers rarely raise an eyebrow at databases that, just a few years ago, they considered impossibly large.

Fortunately, the price/performance ratios of computers and data storage improve by a factor of two approximately every 18 months. However, the demands we place on them often grow even more rapidly. Consequently, the need for database management systems capable of accessing and manipulating very large volumes of data very quickly is even greater today than it was in the past.

The high-performance features of Model 204, along with its ability to manage an almost limitless amount of data, make Model 204 an excellent solution to meet the needs of very large databases and data warehouses.

VERY HIGH VOLUME OLTP AND OLAP

As you automate more of your business processes and provide higher levels of computer support for those processes that retain a manual component, the query and update activity against your databases increases accordingly. The result is more people needing more data from more systems.

However, not all demands are equal. On-line transaction processing (OLTP) generally serves day-to-day operations that must happen quickly – business depends on it. OLTP requests typically access and update just a small amount of data at a time, but, depending on the organization, hundreds, thousands, or even millions of times a day. For OLTP, data update speeds are critical.

In contrast, On-line analytical processing (OLAP) often analyzes vast quantities of data at a time. An OLAP request might, for example, analyze the correlation between household income and purchasing habits across an entire multi-million record customer and prospect database.

Furthermore, OLAP sessions are often fishing expeditions – looking for answers when you do not yet know the right question. Therefore, databases that support OLAP requests must offer very high speed indices on a number of columns to serve the ad hoc nature of the requests.

Model 204 serves both these needs simultaneously. Its multiple advanced indexing technologies ensure the fastest possible access for complex OLAP queries, even against very large databases. Yet, the technology also ensures very high speed updating of data and indices to meet the needs of some of the highest volume operational systems in the world.

MAINFRAME DATA SERVERS

When businesses first adopted the client/server architecture, they saw it as an alternative to the mainframe. Using client/server computing, they distributed data among many servers throughout the enterprise, often outside the information technology department's "glass house."

However, this model presented a number of problems:

- ▲ The cost of managing many distributed servers is typically considerably higher than the cost of managing the one mainframe they replaced.
- ▲ Software costs may increase since each server requires its own set of software licenses.
- ▲ Security becomes more difficult to implement and enforce when servers leave the data center.
- ▲ The new nontechnical managers of the data servers often do not exhibit the same rigor in the critical task of backing up data.
- ▲ The midrange computers used as servers do not offer the same overall throughput of mainframe computers. Therefore, a data server supporting a particularly heavily used enterprise database may become overloaded.

Because of these issues, a number of companies have reevaluated their client/server decisions. While a few see a full-scale return to host-based computing as the answer, the industry as a whole sees a new role for the mainframe: a super server for intranet and extranet applications. By bringing critical data back under the management of a central server, the information technology department can reduce costs, better manage the data, and facilitate enterprise-wide business information delivery.

As enterprises regroup crucial corporate data on to a mainframe server, the demands on mainframe databases increase. The Model 204 profile of very high-performance data queries and updates, support for very large databases, and superior client/server and Internet connectivity through WebGate, make it a technology leader in this mainframe database renaissance.

DISTRIBUTED APPLICATIONS

Even with the trend toward regrouping of some critical data on to a centralized mainframe database, today's computing environments still include a broad range of technologies – mainframes, UNIX and NT departmental and Web servers, AS/400 systems and PCs. This diversity introduces a risk of creating stranded islands of information due to the complexity of the interconnections. However, to maximize value for the enterprise, you must be able to access and combine information from across these platforms easily and quickly, regardless of location or technology.

You can bridge your disparate technology islands using Horizon, MQ/204 and Connect★. This open distributed computing environment allows you to put each computing platform to its best use, but seamlessly share data as required.

In addition, since Horizon and MQ/204 are location-transparent, you can develop and test distributed applications on one machine and install them on different production systems. Because external references are symbolic rather than a physical location, you do not have to change any program logic as you move applications or data around your enterprise.

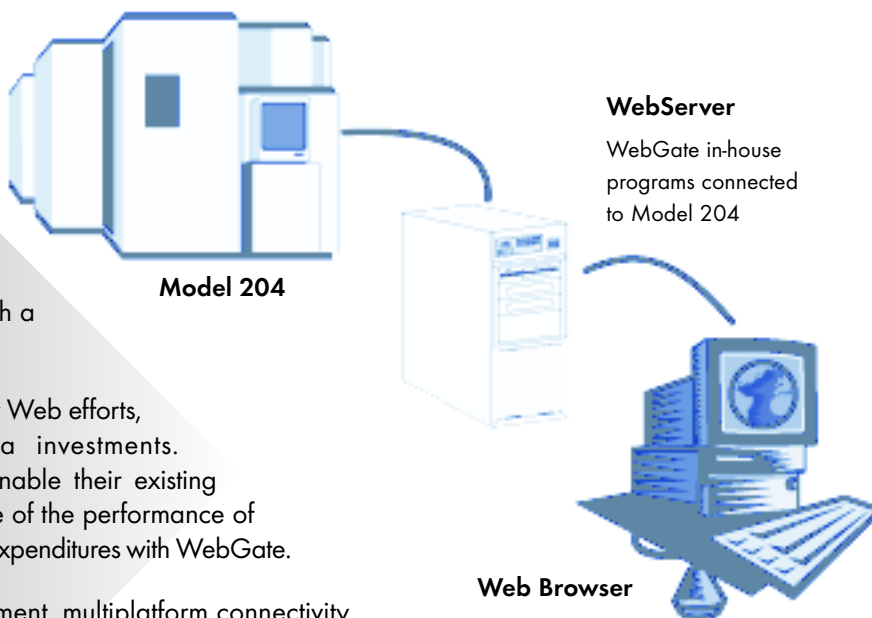
THE WEB

Effective use of the Web can increase revenues by expanding your market reach and offering new products and services. At the same time, it can serve to lower costs through online delivery of information that you previously provided through a physical medium.

However, to gain maximum benefit from your Web efforts, you must leverage your existing data investments. Model 204 customers can quickly Web-enable their existing data and applications to take full advantage of the performance of Model 204 while minimizing any additional expenditures with WebGate.

WebGate includes a development environment, multiplatform connectivity capabilities, and objects specifically provided for easy access to Model 204 data and transactions using SQL, User Language, or both. You then write scripts or programs in the PC-based language of your choice that run on a Windows NT server.

By providing Web browser access to Model 204 data, you automatically give your applications user interfaces with near universal portability. Furthermore, what works for the Internet also works for intranets and extranets. Therefore, you can quickly provide staff, customers, suppliers, and business partners with access to Model 204 data and applications, without individual considerations of their computer hardware or software and networking platforms.



SUMMARY

Computer Corporation of America has always been known for technological innovation in the database management arena. Coupled with a rich set of fully integrated, technologically advanced features and functions, Model 204's ever increasing performance, capacity, and scalability provide unparalleled capabilities to optimize business benefit.

Model 204 solutions provide businesses, governments, military, and nonprofit organizations with low cost of ownership for both the development and ongoing operation and maintenance of applications, with the confidence that they have an ever evolving technology that will allow them to plan for growth.

Organizations benefit from their use of the Model 204 product family in a number of ways:

- ▲ The highly productive rapid, iterative, application development approach that Model 204, User Language, and Workshop/204 make possible combines prototyping, refinement, and production in a single unified process, speeding application development. Consequently, you can deliver higher quality and lower cost solutions sooner.
- ▲ A commitment to standards ensures open solutions and a protection of your data investment even in times of rapidly changing technologies. The SQL interfaces, along with APPC (LU 6.2), TCP/IP, MQSeries, and other connectivity options provide exceptional interoperability even in sophisticated and disparate distributed computing environments - including the Internet, intranets, and extranets.
- ▲ The comprehensive security offered in Model 204, Horizon, and through interfaces to popular security software ensure the integrity of your data.
- ▲ Support for complex data types, a flexible data model, as well as fluid databases and applications, ensure a level of flexibility that allows your databases and applications to evolve in step with your changing business conditions.
- ▲ Model 204 delivers proven, unparalleled performance for relational database management. Customers around the world use it to support exceptionally large databases, very high transaction volumes, thousands of concurrent users, and mixed complex query and update applications. In these most exacting situations, Model 204 users often see performance gains of more than 100 times over other solutions. In some customer benchmark studies, other databases used several times the amount of computer resources and provided response times that were more than 1,000 times longer than Model 204.
- ▲ Performance enhancing options such as Multiprocessor/204 and Parallel Query Option mean that your Model 204 solutions can scale up to meet your requirements no matter how large they may grow.

Years of experience prove the strength of the Computer Corporation of America's database technology. Ongoing product support and development ensure that the technology continues to lead the way in providing business value. High performance, capacity, flexibility, security, openness, and application development productivity make the Model 204 product family the choice of the most demanding organizations in the world.



P.O. Box 9378 ▲ 500 Old Connecticut Path

Framingham, MA 01701-9378

508.270.6666 ▲ 508.270.6688 fax

<http://www.cca-int.com>

Model 204, Imagine and Imagine/CUA are registered trademarks of Computer Corporation of America. User Language, Workshop/204, WebGate, Horizon, Connect★, Sage/204, and Access/204 are trademarks of Computer Corporation of America. All other trademarks and tradenames are used to identify entities claiming the marks and names of their products.

Rev: 3/98