



An Overview of the IBM Tivoli Common Reporting Initiative

This article outlines the Common Reporting initiative and how Tivoli plans to make significant improvements in reporting of high-value information. (Nov 2007)

From a forthcoming article to be published on www.ibm.com/developerWorks



IBM Tivoli is delivering the **Tivoli Common Reporting** software, aiming to provide customers with a reporting solution that is shared across the Tivoli portfolio and can be used by Tivoli, business partners, and customers as a base for increasing the value of solutions built around Tivoli products.

As Tivoli has grown in the number of its products, both internally developed and acquired, the amount of data we collect and analyze has grown significantly. Coordinating all of that data and turning it into useful information has become more and more of a challenge, but an important one for us to address. As a result, we are in the midst of a major initiative to deliver new and exciting technology to make this possible. The Tivoli Common Reporting initiative is more than a new tool: it's a complete reevaluation of how information should be represented, delivered, and shared effectively.

Goals of the Common Reporting initiative

Converge around an incredibly simple launch point for reports, where customers can see what's available and use it with no training.
 Provide an effective way to deliver and share reports and reporting ideas, and find the set of content that gives the greatest value.
 Focus the individual reports on quality of information and value, not on the quantity of reports available.
 Provide consistency of the content and enable linkage between reports for best analysis flow, even if they cross individual product boundaries.
 Integrate reporting and dashboard displays so customers see consistent information and appearance, and can move easily between them.

An effective approach requires focus on both the tooling infrastructure and the content of the reports themselves. We have major goals in these areas, which we believe will have a significant impact on a customer's ability to drive process improvements.

First, we are providing a simple launch point for reports, similar to the way millions of people can deal with media and entertainment, such as music or video clips. Our intent is to have even the casual user community able to find the reports they need and use them with no training.

Second, we believe it is critical for customers to have an effective delivery system for obtaining the right reports, sharing them, providing feedback on how they are used or might be improved, and obtaining them from a wide variety of sources. We are building an open industry around reports that promotes best practices, new analysis ideas, and new visualization techniques. We are already seeing this occur in each of these areas and have every reason to believe that a social network of reporting will continue to evolve and improve this value.

Our third goal is to focus on the individual reports that we deliver with our products, or even across the products. We are spending more time with customers to determine what information is critical to their business processes, and building that back into the content of reports. In the process, we are spending more time evaluating how we visually represent that content so it is clear and makes the greatest impact as quickly as possible. We believe the best reports are those that draw your interest and allow you to make numerous observations from the visuals.

Fourth, we want you to easily understand and compare that information across all the reports we provide. You should see dates and times displayed in the same way; data fields should be consistently formatted and have the same labels when they appear in multiple places. The metrics and scales should be comparable between reports so they don't lead to incorrect assumptions. The colors and fonts should be configurable and consistent.

Finally, we realize that information content is what's important, not the tool that displays it. Putting the information into a report for printing or distribution should provide the same information available to the operation's enterprise console or executive dashboards. There should be consistency of that information, even though one may be viewed in near real-time and the other distributed as a snapshot once a week.



The Tivoli Common Reporting architecture...open and simple.

In order to satisfy those goals, we began with an architecture that satisfies different product requirements and different customer requirements, while keeping it as simple as possible. It is also obvious to us that report authorship is a creative and personal process, and that there are many, many sources of good reports. It is important to make our solution open to encourage creativity and produce a report marketplace. We chose the Business Intelligence Report Tool (BIRT) technology from the Eclipse open source project, and its XML-based report definition tags; and we built a set of services to make a complete solution. The definition can be easily provided by third parties, and has a growing list of companies already providing compatible content. The BIRT report authoring tool gives anyone the opportunity to create and customize reports. The following diagram shows the overall architecture for the solution:

Tivoli Common Reporting Architecture

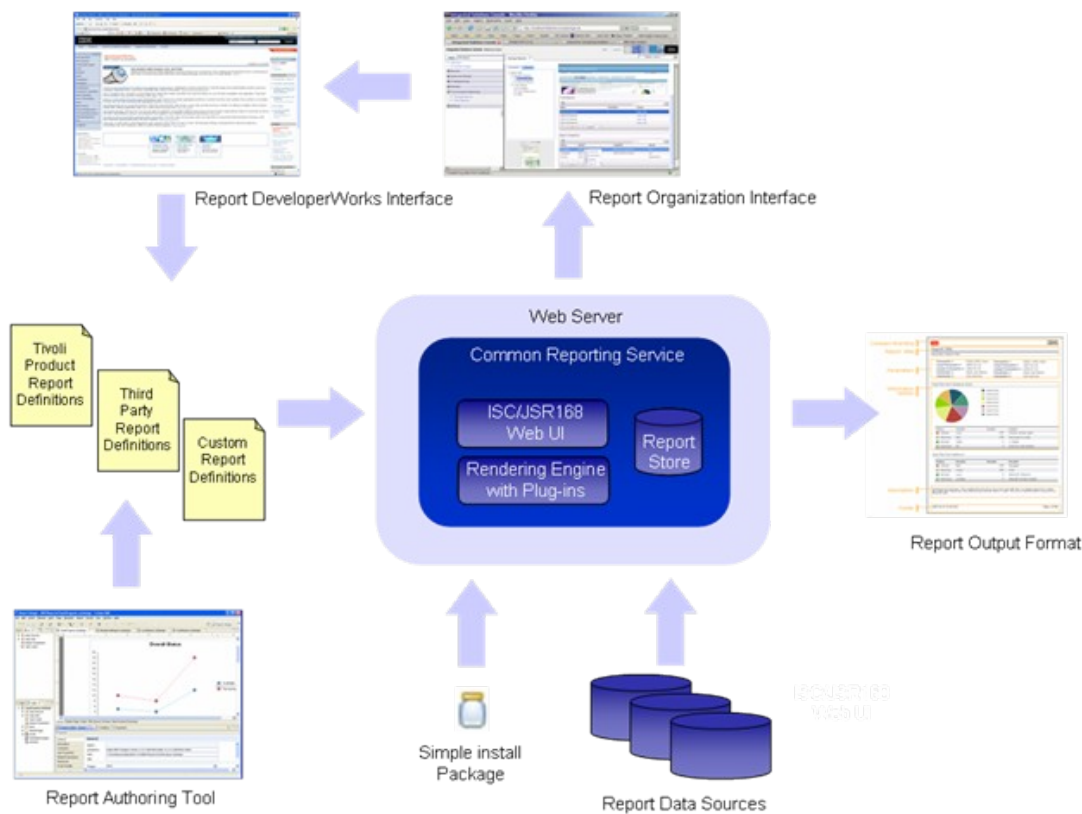


Figure 1. Tivoli Common Reporting Architecture

It begins, of course, with a simple install package that lays down the necessary infrastructure and services. In the bottom left, the authoring tools are used by the report creator, which results in a report definition, shown in the picture. This definition has standard XML tags describing every aspect of the report, from its data source to its appearance. This file might also point to additional resource files, such as libraries and images. These files are all compressed together as an independent deliverable, which can be delivered in a variety of ways, and then imported into the Tivoli Common Reporting data store. Once there, the report navigation interface displays the list of reports for execution or maintenance. When a report is executed, parameters can be provided; then the system pulls data from the configured data sources and produces the formatted report. There are different plug-ins to create different output formats, including HTML, PDF, and Microsoft Excel™ files.

At the top left, the Tivoli Common Reporting site at IBM developerWorks provides the ecosystem for delivery and download of new reports from different sources. It also forms a community for sharing ideas and information about reports and their content. It is an important part of the user experience and is integrated into the report navigation interface.

The report navigation interface...your portal to information.

There must be a starting point for finding and organizing the available reports, launching them, scheduling them, distributing them, and controlling access to them. The Tivoli Common Reporting service is a web-based application to do all that, and do it simply and efficiently. The notion of organizing reports isn't much different than organizing your media files, such as music, video, podcasts, or pictures. The main difference is that reports can be slightly different each time you run them, based on the updates to the underlying data. As a result, each report can have different **snapshots**, each of which represents a particular set of parameters and the particular time that it was executed. You can create a new snapshot any time you wish, and you can keep the snapshots around for later reference or distribution.

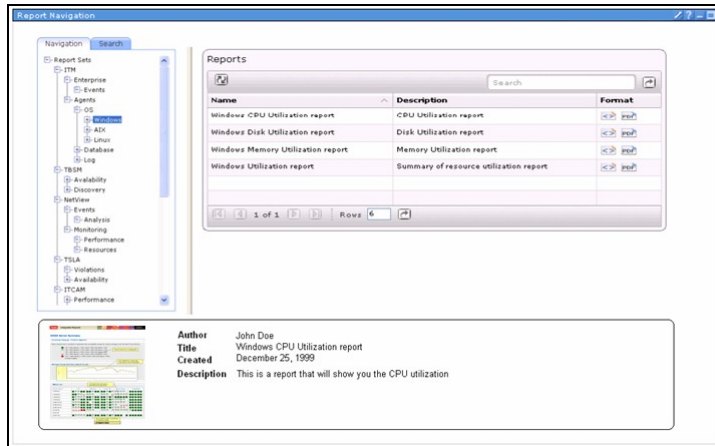


Figure 2. Concept of the Main Report Organization View

the Tivoli Common Reporting service. A customer can choose to configure this to point to their own information instead, showing important company news for example.

Finding and organizing reports is done with **report sets**, which are fixed folders or “playlists,” and in the future, we’ll provide “smart sets” whose contents are determined by specified criteria (for example, the most recently executed reports, or the ones used most often). It’s also possible to search across the report titles, keywords, or other information about the report definition.

The Tivoli Common Reporting developerWorks Interface... the ecosystem for report creativity.

It’s one thing to provide a tool for viewing reports; it’s an entirely different matter to provide a complete set of reports that meet the specific needs of the large number of customers that will use them. While we are greatly improving the quality of our out-of-the-box reports using our experience with IT process best practices, customers will also want to customize those reports or provide entirely new ones. There is a great opportunity to expand the reporting value of the data our products gather by growing an open report marketplace. This marketplace already exists, and we believe this strategy will accelerate its growth, giving customers more options and more value.

“It was important to make our solution open to encourage creativity and produce a report marketplace.”

The Tivoli Common Reporting site at IBM developerWorks is the place where all this shared information can be obtained. This Web site is part of IBM’s support for the development community, with tools and assistance in many forms. The site provides articles, references, new information, opportunities to communicate with the development and architecture team, along with downloads of the service, tools, and new reports. It also provides the ability for customers and partners to share report ideas and sample reports for download, and to post comments. We anticipate everything from “how to use the tools” to “statistical process control techniques for reporting.”

We are incredibly excited about the possibilities for creating this ecosystem and watching it grow across the customer and provider community, moving the industry forward in analysis and reporting techniques, and seeing the resulting improvement in the value our customers get from their investment in Tivoli products.

This is also the location to search for reports that are available in between product releases or that combine information across products. You can see reports that provide information you need from a product you might want to purchase. We will also have report samples and ideas, with the opportunity for customers to provide feedback and even vote for the visual



approaches they like best.

Report content and output...*the real value and impact to the customer.*

It's one thing to provide a great tool set for finding and launching reports, but it's even more important for those reports to provide insightful information and analysis that a customer can use. We have put considerable focus on the content that is being delivered in those reports, concentrating on the following areas:

Getting the right data...

- ✓ More analysis of the goals for each report, and what actions can be done with the results
- ✓ More work with product teams to get the underlying data stores consistent, with values that allow linkages across them

Getting the right visual display...

- ✓ Better report effectiveness and "information density"
- ✓ New and improved graphics capabilities

Getting consistent standards and guidelines...

1. Standard report layouts, for better comparison across reports
2. Standard color and font schemes
3. Standard data representation, units, and labels
4. Standard parameterization and linkages
5. Standard report definitions (common libraries, templates, etc.)

The first part of the list is all about the development of useful information from the data sources, with a clear and concise goal in mind. If a report (or set of reports) is designed to help optimize server performance, for example, then it should provide the information necessary for the administrator to make the exact adjustment required for that optimization. In addition, we are working with each product team to drive requirements back into the data stores for improvements in data quality, consistency, and added value to create new reports.

The second part deals with the best visual display for the desired effect. Doing this takes good analysis on not only the data, but the operational and business processes that use the information and the potential actions that might result. There is some great material on this subject in a set of books by Edward Tufte [1], which covers the notion of **Information Density**, a focus on the amount of information relative to the size of the display. It's also critical that the reports maintain visual integrity, where they aren't deceptive in the presentation of the results.

The reports that our customers deal with are typically used for gaining insight from IT operations processes to see if some action needs to occur. So producing the right reports is important to help develop **continuous process improvement** in one of the three areas in the diagram below.

The current state of the art in the industry is to simply monitor IT resource data, and show historical usage. There might be trending provided, but it usually deals with a single element. There is very little real analysis/conclusion, cause/effect, or other relational comparison reports which help to draw conclusions more easily. There are also very few reports that relate the resource usage or availability to the actual business cost drivers that we outline in the diagram.

Improvement of Cost	<p>The focus is on providing <i>optimum service at minimum expense</i>. It may be results per headcount or overall cost reduction. This generally includes:</p> <ul style="list-style-type: none"> ✓ Basic productivity gains in shortening process event times ✓ Ability to improve throughput for business service transactions ✓ Process capacity increase due to less errors and rework
Improvement of Quality	<p>The focus is typically on the process <i>yield</i>, or the amount of results that complete the process correctly compared to the amount started. The following are all critical to quality:</p> <ul style="list-style-type: none"> ✓ Accuracy of the process artifacts (for example, a change implemented on the correct resources) ✓ Intended results are obtained without incident and rework ✓ The process's ability to delivery continuous service
Improvement of Schedule	<p>The focus is typically the reduction of process steps, the process cycle time, and the ability to accommodate <i>adjustment in velocity</i> when business conditions warrant. Schedule factors are:</p> <ul style="list-style-type: none"> ✓ Flexibility for expediting certain process occurrences ✓ Optimization of individual process steps ✓ Optimization in the coordination of multiple occurrences ✓ Predictability in the process for estimation and management

A major part of this initiative is to provide report content that helps drive process improvement: cost improvement, quality improvement, and schedule improvement. We will provide reports that show how our own management products help the IT operations staff make these improvements.

We have plans for unique ways of bundling reports, so they can be used like cookbooks for performing well-known analysis procedures. For example, a set of reports might take you through a series of steps to analyze server performance for a typical Web server:



Instead of 500 individual reports, they are bundled



into 50 standard analysis packs that are the most-used scenarios based on our experience with customer feedback.

The third part of the list deals with standards and guidelines for all reports to follow. These standards provide consistency in the way customers see the reports, regardless of where they were created or what data sources they present. This

Figure 3. Process Improvement Factors

consistency goes a long way towards providing the visual integrity mentioned previously, which deals with providing the best possible way of communicating information to the people that need to understand it.

There are several aspects of consistency, as you can see from this list. We have produced an IBM Tivoli Common Reporting *Development and Style Guide* that goes into considerable detail on each of these areas. Our own products are reviewed and will comply with these standards, and the guide will be available for partners and customers to follow as well.

Each report has a standard layout of sections that include a top configurable graphic splash, a list of parameter values used to generate the data values, a report description, and footers. There is a default color and font scheme, with the possibility in the future of different configurable “skins” to allow customer preferences. All of these are shared through report libraries and templates, which allow all reports to share a single part of the definition, rather than repeat the redundant information in each. This is also true of the configurable data source locations, which make installation easier when the customer points to the unique location of their reporting databases.

Another critical aspect of consistency is the actual data syntax, which covers everything from date/time syntax to similar units of measure. When you see historical charts, for example, you want to assume consistency in both the scaling as well as the rate, so you can compare across reports accurately. We are also working to ensure that data labels are the same across the report set, so that, for example, a data value called “Memory Usage” refers to the amount of virtual memory used, not the amount of physical memory.

Finally, we have standards and best practices for how parameters should be defined, enabling hyperlink possibilities between reports. For example, any reports dealing with the aspects of a server should always provide a parameter for the hostname for either filtering or grouping.

Report authoring tools...unleashing creativity and customization.

Even with all of our work to create industry-leading report content, we recognize the incredible demand for reports that meet very specific customer needs. Process improvement metrics are far from standardized, and each customer has specific IT configurations that can have subtle differences. The ability to both create new reports as well as modify the reports we provide must be available, and must be easy to use. Customers have been very clear about the need to push a reporting tool out to a wider audience that does not have the background skill to use an engineering tool for report creation.

We have decided to take a multi-pronged approach to authoring and customization, designed to reach users with different skill sets and provide a wider set of options when required. We’ve categorized those approaches as follows:

Simple Customization of Standard Reports: A customer uses the report navigation interface to make appearance changes in an existing report, without really changing the basic structure or content. Obviously, filtering parameters are provided this way, along with options like title graphics and text, sort order, including certain attribute values, changing their labels, and font or color choices. These are relatively easy to understand and change without any training, intended for a wide audience.

Simple Report Generation Wizard: A customer selects a basic report structure from a set of thumbnail examples, then select from a set of data content templates to fill the report, and finalize any appearance items. This approach has limitations on the amount of customization initially done, but is also intended for a wide audience to let them become self-sufficient for getting basic information. This will also be a part of the report navigation Interface.

Dynamic Report Combination: A very interesting and rapidly growing area is the use of Web 2.0 technologies to combine application parts into new solutions quickly and easily. Reports are a perfect fit for the “mashup” and widget approaches becoming popular. We can provide tools and interfaces enabling customers to link reports together within their own application flows for customizing a solution.

Advanced Report Creation and Customization: Inevitably, there are some reports that are sufficiently complex that they require a more skilled creator and a more power creation tool. For these reports, we have the Eclipse BIRT report designer authoring tool which can meet these more advanced requirements. Even this tool can be used to quickly create simple reports as well, as its popularity in the open source arena has shown.

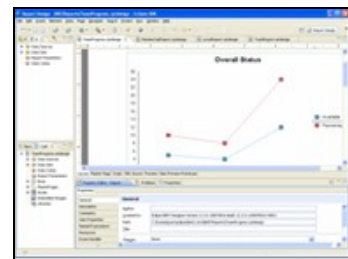


Figure 4. Eclipse BIRT Designer

The first three of these are not available yet, but we are working with customers to determine which techniques work best for them. Each of these approaches has an important part to play, and we believe customers will find them integrated together in an easy to understand fashion.



The Common Reporting software...*the power behind the scenes.*

Notice that we've shown you all the interfaces from the customer's point of view, giving you an idea of the capabilities and simplicity of the strategy. We wanted to make sure you understood that we are focused on the customer experience first and foremost. We've intentionally left the infrastructure for last, since IBM is known for building solid, dependable, and high-performance systems. The Tivoli Common Reporting software is no exception, using our reliable and business-capable middleware for its runtime services.

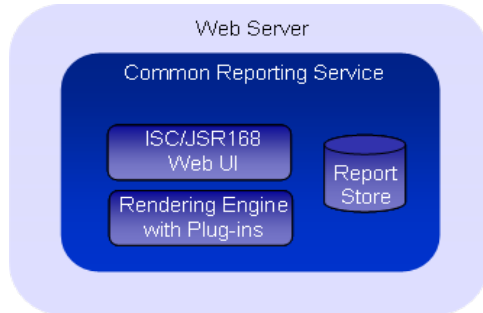


Figure 5. Reporting Service Components

The report data store is an embedded database system coordinated with the file system to provide high performance search and data management with no advanced administration requirements. Since it uses standard SQL and JDBC connectivity, it will allow customers to use their DB2 or Oracle installations in the future if desired. These options will give us large-scale capability and integration into backup and recovery policies at production sites.

The BIRT open source rendering engine is included within the Tivoli Common Reporting software, which runs inside a small-footprint WebSphere application server known as eWAS, again making the install package and configuration simpler. In the future, this can easily be configured to run in a full WebSphere production environment with failover and load balancing. The Web UI uses portlet technology, for flexibility in customizing the user experience with the web browser.

This strong combination of lightweight middleware with the option to move up to the full versions in the future gives us a very simple installation environment and the ability to scale up to whatever customer requirements dictate.

There are a few other design features we believe customers will appreciate. Every function the user interfaces provide is performed through an available API, and there is a command-line interface (CLI) as well. These interfaces allow for custom solution scripting and application integration whether done by the customer or a service engagement. These can be used for operational maintenance and operation or as part of a work flow within other IBM Service Management processes.

The service will provide some instrumentation interfaces and the agent code for automatically monitoring its operation using our own Tivoli products, and alerting when problems arise. We also plan to provide some reports about our own service as part of the installation in an upcoming release.



Tivoli's Common Reporting initiative...wrapping it all up.

As you can see, we're quite excited about where we are going with our reporting initiative, from the advances in tooling, to the ecosystem of information and feedback, to the value of the content itself. Tivoli is making a significant investment to provide customers with the best possible solution for gaining insight on their IT operations. Despite this investment, the reporting tools, services, and content are all available at no charge. They come as part of the products that you purchase.

"Here's the best part...we're providing all these tools, services, and content free of charge."

We think you'll agree that this is an important step forward for Tivoli and their customers. It is important to us that we get your feedback on the tools, the individual report value, the ways we can improve the solutions. The Tivoli Common Reporting site at IBM developerWorks is the place for that, with blogs and forums, and we look forward to hearing from you.



Ron Baker

Senior Technical Staff Member, IBM Tivoli Software
rbbaker@us.ibm.com

[1] *The Visual Display of Quantitative Information*, Edward R. Tufte, Graphics Press, Cheshire, Connecticut, 1997. See also, *Visual Explanations*, Edward R. Tufte, Graphics Press, Cheshire, Connecticut, 1997.